Library and Archives Canada

Published Heritage Branch

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque et Archives Canada

Direction du Patrimoine de l'édition

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence
ISBN: 978-0-494-29994-4
Our file Notre référence
ISBN: 978-0-494-29994-4

NOTICE:
The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:
L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

Canada

Abstract

Economic feasibility determines if mine projects make it beyond the design stage. The design process is a major component. Designs may be updated several times to achieve minimal operating costs and maximal profit. This process is a form of optimization. An updated optimal design can be significantly better than a design that is arrived at quickly.

This thesis covers two areas of underground mine design that are amenable to optimization: the design of individual stopes and the sequence that the stopes are mined. The problem of stope geometry optimization is developed with an optimization algorithm. A framework for stope sequence optimization is then developed. Both areas of design optimization are applied in case studies. Stope geometry optimization showed good economic improvement ranging from 14 to 33 %. The percentage improvement achieved by stope sequence optimization was a modest 1 to 4 %. A framework has been established for further development.

Acknowledgements

I am grateful to Clayton Deutsch and Oy Leuangthong for their invaluable support and encouragement they gave during my degree. Their friendship and professional relationship is very important to me.

I would also like to thank Placer Dome, Barrick Gold and Tech Cominco as well as other industry affiliates of the Center for Computational Geostatistics (CCG) for initiating ideas and providing guidance in this work. Financial support provided by Placer Dome, the Natural Sciences and Engineering Research Council (NSERC), and the University of Alberta is greatly appreciated. My gratitude is also extended to fellow colleagues at CCG for their assistance.

Finally, I would like to thank my fiancée Erin. Her patience and support were invaluable during this research.

Table of Contents

Chapter 1 Introduction 1
1.1 Background 2
1.1.1 Geologic Complexity 2
1.1.2 Numerical Modeling 3
1.1.3 Mining Methods 5

Chapter 2 Numerical Optimization 9
2.1 Optimization with Continuous Parameters 10
2.1.1 Simulated Annealing 10
2.1.2 Genetic Algorithms 11
2.2 Optimization with Discrete Parameters 13
2.2.1 Exchange Heuristics 13
2.2.2 Branch-and-Bound 14
2.3 Techniques Applied 15

Chapter 3 Stope Geometry Optimization 16
3.1 Parameterizing a Stope 16
3.1.1 Stope Geometry 16
3.1.2 Constraints 17
3.2 Block Model Clipping 22
3.2.1 Primary Clipping 23
3.2.2 Secondary Clipping 24
3.2.3 Tertiary Clipping 24
3.3 Stope Modification 25
3.4 Optimization with Simulated Annealing 26

Chapter 4 A Framework for Stope Sequence Optimization 28
4.1 Parameterizing a Sequence 29
4.1.1 Stope Parameterization 29
4.1.2 Development Openings 30
4.1.3 Parameterizing Operations 30
4.1.4 Constraints 33
4.2 Choosing a Feasible Sequence 34
4.3 Stope Sequence Scenario 36
4.4 Optimization 38
4.4.1 Optimization with Simulated Annealing 40
4.4.2 Optimization with Probabilistic Decision Making 42

Chapter 5 Case Studies **46**
5.1 Application: Stope Geometry Optimization 46
5.1.1 Case 1: The Objective Surface 48
5.1.2 Case 2: Sublevel Stoping 51
5.1.3 Case 3: Cut-and-Fill 53
5.2 Application: Stope Sequence Optimization 55

Chapter 6 Conclusions and Future Work **61**
6.1 Stope Geometry Optimization 61
6.2 Stope Sequence Optimization 62
6.3 Future Work 63

References **65**

List of Tables

3-1: Potential problems, parameters and analysis for underground hard rock mine openings. 20

5-1: Sublevel stope optimization summary. 52
5-2: Cut-and-fill stope optimization summary. 54
5-3: Stope sequence optimization results for simulated annealing and PDM. 57
5-4: SA NPV Statistics 60

List of Figures

1-1: Flowchart describing a deposit modeling process. 4

1-2: Schematic of sublevel stope drilling techniques. 6

1-3: Schematic of the vertical crater retreat mining process. 7

1-4: Schematic of raise mining. 7

1-5: Schematic of a level plan for drift-and-fill mining. 8

2-1: A function and a local and global minimum. 9

2-2: Flowchart showing a simulated annealing algorithm. 10

2-3: General flow of a genetic algorithm. 12

2-4: Processes of crossover and mutation. 12

2-5: Schematic of a set of stopes, their sequence and the associated state-space tree. 14

3-1: Schematic showing stope geometry components. 17

3-2: Components of stope geometry that can be constrained. 17

3-3: Schematic of minimum mining width constraint as applied to cut-and-fill mining. 18

3-4: Schematic showing a stope design leading to ore hang-up problems. 19

3-5: Schematic of maximum overall face dimensions given constrained triangulated face dimensions. 20

3-6: Schematic of stope cross sections showing face deviation limits imposed by drilling equipment. 21

3-7: Schematic of corner deviation limits imposed by drilling and blasting. 21

3-8: Comparison schematic of a stope object with a view frustum. 22

3-9: Schematic of a regular grid and its components. 23

3-10: Schematic of block status as determined using a z-buffer. 24

3-11: Schematic of block refinement process. 25

3-12: Flowchart of the simulated annealing algorithm applied to stope geometry optimization. 26

4-1: Schematic of an ore body, panel and stopes. 28

4-2: Jackleg drill and drill jumbo. 31

4-3: Schematic of a procedure for advancing along development openings. 32

4-4: A scoop tram and mine truck. 32

4-5: Schematic of basic panel extraction procedure. 34

4-6: Schematic of a mining scenario indicating which stopes are infeasible given constraints. 35

4-7: Schematic of a partially completed sequence scenario for state-space tree analysis. 35

4-8: A state-space tree for the sequence scenario from Figure 4-7. 36

4-9: Schematic of the simple stope sequence scenario. 36
4-10: Schematic of a panel segmented into individual jobs. 38
4-11: Data organization for a two dimensional panel of stopes. 39
4-12: Simulated annealing algorithm for stope sequence optimization. 41
4-13: Process of accepting suboptimal sequences to reach more optimal solutions. 42
4-14: Flow chart of the PDM algorithm. 44

5-1: Three cases for stope geometry optimization. 47
5-2: Histogram of block values for a synthetic economic model. 47
5-3: Sample block model displaying blocks with positive value. 48
5-4: Case 1 stope, block model, and blocks for value calculation. 48
5-5: Various levels of complexity of stope value surfaces for case 1. 50
5-6: Plan views of the three levels used in designing the sublevel stope. 51
5-7: Sublevel stope vertex constraints. 52
5-8: Initial and optimized sublevel stope designs. 53
5-9: Plan view showing cut-and-fill drift design. 54
5-10: Intersection of the block model with the cut-and-fill stope by full blocks and refined blocks. 55
5-11: Initial design and optimized design for the cut-and-fill stope. 55
5-12: Panel of 70 stopes used for the stope sequencing application. 56
5-13: Two-dimensional projection of the sublevel stoping panel. 57
5-14: Optimization progress chart for simulated annealing and PDM. 58
5-15: Stope sequences for simulated annealing, PDM and the difference between them. 59
5-16: Average and standard deviation of SA sequences. 60

List of Symbols

x,y,z	Location in Cartesian coordinates
x_0,y_0,z_0	Origin of a regular grid specification
dx,dy,dz	Dimensions of a regular grid block
x_{index}	Regular grid block x coordinate index (same for y and z)
w	Width
h	Height
r_H	Hydraulic radius
F_0	Initial value of an objective function
F_k	Value of an objective function at state k
T	System temperature for simulated annealing
V	Volume of a regular grid block
P	Price of a commodity
$Prob$	Probability
R	Recovery of a commodity from mineral processing
C_m	Cost to mine ore
C_p	Cost to process ore from mining
$\mathbf{u}$	Cartesian vector
$I(\mathbf{u})$	Value of a regular grid block centered at location $\mathbf{u}$
b	A particular grid block from a regular grid
f_b	Fraction of grid block b inside a stope
v_b	Value of grid block b
S	A single stope object or matrix of stopes
D	A single drift object or matrix of drifts
SP	Matrix of stope profits for a panel
DP	Matrix of drift profits for a panel
ST	Matrix of times corresponding to mining stopes in a panel
DT	Matrix of times corresponding to mining drifts in a panel
b	A set of regular grid blocks within a stope
t_a	Time to access a stope in a panel

t_p	Time to prepare a stope for mining
t_e	Time to mine a stope
t_{dm}	Time to mine a drift
δ	Logical indicator
r	Annual discount rate
NPV	Net present value of a panel of stopes
$\mathbf{P}$	Matrix of probabilities for a panel
$\mathbf{F}$	Matrix of stope properties
$\mathbf{w}$	Matrix of weights applied to stope properties
$\boldsymbol{x}$	Vector of parameters for an n-dimensional function
$f(x)$	A function evaluated with parameter x
$\frac{\partial f}{\partial x}$	Partial derivative of a function f with respect to x

Chapter 1 Introduction

Evaluating and improving the economic potential of a mining operation is a primary objective of mining engineers. This is true for both surface and underground mines; however, most currently available tools are aimed at open pit operations. Technological advances made in mining equipment and practices have allowed underground mines to descend to immense depths. Both increasing depth and lower grades are driving mines towards a more marginal state.

With underground operations becoming more marginal, tools for evaluating and enhancing economics have become more popular. Two particular areas of interest in this regard are stope geometry and stope sequence optimization. The problem of stope geometry optimization involves taking an initial stope design and maximizing its economic return by modifying the design. Underlying information to accomplish this comes from geological, geotechnical and grade models along with economic and mining method-specific parameters. Optimization of a stope sequence involves choosing an order to mine stopes, which also requires scheduling of equipment crews and other mine operations, to maximize the present value of that set of stopes. These optimization tasks can be complex problems to solve.

Each mine has its own constraints and site specific considerations. Mine design and operation depend on structural, geological, geotechnical and mineral grade properties of the deposit, each containing some level of uncertainty. Geologic models describe various minerals, host rock types, jointing and fault plane occurrences. These features are used in other numerical models including geotechnical and grade models. From a geostatistical perspective, mineral grade is considered a random variable and is modeled using simulation techniques. This may also be true for geologic features such as mineral and host rock types. Uncertainty is expressed through a set of possible geological settings and mineral grade distributions.

The variety of available mining methods makes the problem of optimization more difficult. Classic mining methods include sublevel stoping, vertical crater retreat and cut-and-fill stoping. Each of these methods has a number of variants: raise mining is a form a sublevel stoping; cut-and-fill can be underhand of overhand. Mines may also employ a combination of methods in different locations. Every mine will have different geometry and mining constraints. Constraints are important to the optimization process as they prevent a stope from being modified into something unsafe or impractical to mine. They control how a stope is optimized. A stope designed for sublevel stoping will undergo a different modification process than one designed for cut-and-fill mining. Mining method also influences stope sequence optimization. Each method comes with different development design (drifts, stopes, cross-cuts, ramps, etc...), mine operating procedures and limitations.

A reasonable solution to these problems is achieved with the array of computational optimization algorithms available. Regarding stope geometry optimization, techniques such as simulated annealing (SA) [1] or genetic algorithms [4] can be adapted and applied. Stope sequencing can also be solved using these techniques, but a more appropriate approach may be to utilize combinatorial optimization methods like branch-and-bound [3] or ant colony optimization [7].

This thesis will develop the problems of stope geometry and sequence optimization. Numerical optimization is essential to this thesis and will be discussed at the beginning. A solution to stope geometry optimization will be developed followed by a framework for sequencing. Case studies for stope geometry and sequencing will be presented. Additional background on deposit modeling, geologic complexity and mining methods will be provided first.

1.1 Background

Several stages of the mine design process are carried out prior to performing any economic optimization. An initial mine design must be in place that is appropriate for the structure of the deposit, the geologic and geotechnical characteristics, and mineral grade distributions. Understanding the structure of a deposit is important for identifying coordinate systems and grids for modeling. Rock type and strength, jointing, faulting, hydrological, and other features constrain the mining methods considered feasible to exploit a deposit. With an understanding of geology, modeling of other features including the grade distribution can be carried out. Various geostatistical modeling techniques are available for estimating and simulating both geological and grade variables. They could also be used for geotechnical property modeling. An advantage of using geostatistical techniques is that uncertainty can be quantified and accounted for in optimization.

Knowledge regarding geology, geotechnical and mineral grade along with any perceived constraints on mine openings and procedures is utilized in creating an initial mine design. Several variants on design will be developed and evaluated before an acceptable one is carried forward to the implementation stage. An important step is building an economic model and evaluating potential designs based on it.

1.1.1 Geologic Complexity

Complexity of a mineral resource stems mainly from underlying geology. The geological setting influences many other parameters including rock mass quality, the distribution of mineral grades, mineral processing design, and the mining method of choice. All of these mine design components are used in developing an economic model for evaluating design feasibility. It is essential to develop as accurate a geologic model as possible.

Many geologic features are recorded during exploration and during the actual mining stage. Possible features include rock type, mineralogy, texture, hardness and abrasiveness, folding, faulting and fracturing, hydrogeology, drill rate and bit wear, powder factor, ground support requirements, drainage, and subsidence [9]. These features can be directly recorded from core and rock samples or indirectly recorded from process recovery, wear on drilling equipment and effectiveness of explosives. Combinations of features can further be used to delineate a mine into different geologic zones, each imposing different constraints on mining and requiring different mining practices.

From a geostatistical approach, rock type is a commonly used geologic parameter. Different rock-types originate from different underlying geological processes or from different components of the same geologic process. These are identified during the modeling process. Rock type will be predicted throughout a deposit by stochastic means resulting in multiple realizations.

Sampling is an ongoing process over the mine life. Samples from exploration and delineation drilling and during mining are used to keep the affected models as current as possible. Visual and grab sampling helps geologists and mine engineers make immediate decisions regarding daily operations. An example is deciding orientation of a drift being mined in ore. If ore is noticeably different than waste, the ore can be followed based on visual inspection.

1.1.2 Numerical Modeling

Most resource and reserve models in underground mining are constructed using Cartesian regular grids. All cells in the model are of constant dimension in each direction. Each cell contains attributes estimated through deterministic and/or stochastic methods. The initial step in the modeling process is structural modeling. Understanding the type of deposit being dealt with and how it came to exist and its structural features (stratigraphy, faults, etc...) offer insight into directions of continuity and maximal stress through the deposit. These are important for identifying a coordinate system local to the deposit and grid specification for modeling attributes. The next step involves creating a geologic model. Geotechnical properties may be modeled at this point as well. Mineral grade modeling is carried out last as estimation parameters vary by rock type, mineralogy and other zoning features such as faults. A mine design can be devised with structural, geologic and mineral grade models. These three components can then be used to construct an economic model. Mine design and economics become an iterative process to achieve a design yielding maximum profitability. Figure 1-1 provides a flowchart.

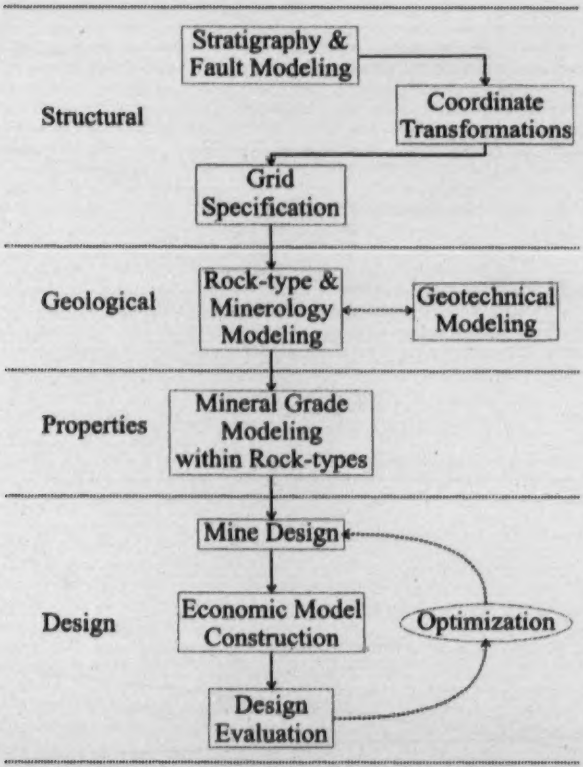

Figure 1-1: Flowchart describing a deposit modeling process.

Building a geologic model is the second step in the modeling process. Any geologic features recorded from sampling and deemed important to mining are modeled to provide a three-dimensional picture of those features. This model is updated as mining progresses and samples are received. Most commonly, rock type is modeled based on mineralogy. Other components of the geology model cover geotechnical features. These play a role in designating zones, constraining mine openings and establishing ground support requirements. Parameters such as intact rock strength, drill core quality, joint spacing, condition and orientation, and groundwater conditions can be measured and used to define the rock mass rating (RMR) designation for example [14]. Faults may also be modeled, possibly as triangulated surfaces. They can be used to guide coordinates when estimating grades.

Estimation of grades is carried out within each rock-type defined from geology modeling. Differences in grade from one rock type to another may involve average grade, variation and spatial anisotropy structure. Utilizing geostatistical techniques to model grades will result in a number of realizations to consider. In regards to mine design, which is discussed in the next section, one cannot consider all realizations because of CPU and professional time. Grade and geology can be summarized in one model as expectations or averages of their respective realization sets. A mine design could be based on these models.

An economic model can be constructed from a mine design and grade model. The economic model is required to evaluate a mine design. As depicted in Figure 1-1, the optimization process would involve making modifications to the design with the intent of increasing profit. Many parameters contribute to the economics of a project and they depend on mining method, equipment and operating procedures, the commodity being mined, the final product produced, and the location of the mine.

For evaluation and optimization purposes, an economic model is summarized by a profit value for each cell. The profit indicates the value of a cell if it is mined and processed to a final product. Parameters involved in determining a cell's value include mining cost (C_m), processing cost (C_p), processing recovery for each product (R_i, i=1,...,n), market price of each product (P_i, i=1,...,n), mineral grades (g_i, i=1,...,n), and cell volume or tonnage (V). In a simplified form, this information can be put together as Equation 1-1. Regarding stope geometry optimization, all cells within a stope are mined and processed, thus their economic value can be evaluated in this fashion. For stope sequence optimization, the cells will be mined at a specified time. Its value must be discounted to evaluate net present value.

$$I(\mathbf{u}) = V\left[\sum_{i=1}^{n} g_i(\mathbf{u}) \cdot P_i \cdot R_i - \left(C_p + C_m\right)\right] \qquad 1\text{-}1$$

Where $I(\mathbf{u})$ is the profit of a cell centered at location $\mathbf{u}$, V is its volume or tonnage, n is the number of commodities, $g_i(\mathbf{u})$, P_i, and R_i are the grade, price and recovery of commodity i, and C_p and C_m are the processing and mining costs respectively.

A profit value can be calculated for all cells, even unprofitable ones. Profit depends on the process used for extraction and the associated operating cost characteristics. Characterizing development operations is important for calculating costs and time to access a mine. Mining access openings including ramps, drifts and cross-cuts are important to the problem of stope sequencing. In the case where development occurs in ore profit is calculated as in Equation 1-1 and discounted depending on the mining schedule.

1.1.3 Mining Methods

A wide variety of mining methods are available for putting mineral deposits into production. Each has different design characteristics and mining procedures. An understanding of these methods is critical to developing flexible stope geometry and stope sequence optimization algorithms. Algorithms applicable to only one mining method and a specific stope design are less desirable than general solutions.

Designing a mine encompasses many decisions, one of the first being which mining method is most applicable given geology, mineralization and economics of a deposit [10]. There are three categories of underground mining methods: unsupported, supported, and caving. One unsupported method is sublevel stoping, which also has several variants

including vertical crater retreat (VCR) and raise mining. Cut-and-fill stoping is a supported method. It can be executed overhand (bottom up) or underhand (top down). All of these methods are ideal for steeply dipping ore bodies that can be mined in a fairly consistent manner.

Sublevel Stoping

Sublevel stoping and VCR are both bulk mining methods that are non-selective whereas raise mining is a more selective variant. Sublevel stoping may also be referred to as blasthole or longhole stoping. These methods are development intensive with most development in ore and are intended for high production rates. All three variants utilize gravity fed ore extraction systems. Applicable ore bodies are large and fairly regular, somewhat strong, competent and dipping steeper than the angle of repose of blasted ore.

There are some important general features of sublevel stoping for design and sequence considerations. A stope remains open while it is being mined, then it is backfilled. Because stopes are typically large, some remain as pillars while others are mined and backfilled. These pillars are extracted when backfill has settled sufficiently. Extraction of pillars may not be possible due to stability issues, in which case ore recovery can drop lower than 80%. High production rates are achieved because it is highly mechanized. An abundance of equipment and large stopes allow drilling in advance and long time gaps between production blasts. Regarding design, stope width in general must be 20 feet or larger for high mechanization. Length and height of the stopes will depend on geotechnical analysis with heights ranging from 60 to 400 feet [9]. Several kinds of drilling can be applied to sublevel stoping including parallel, ring and fan drilling, see Figure 1-2. Parallel drilling offers maximum drilling efficiency for very steep dip, whereas ring and fan drilling are applied to massive ore bodies (large in all dimensions).

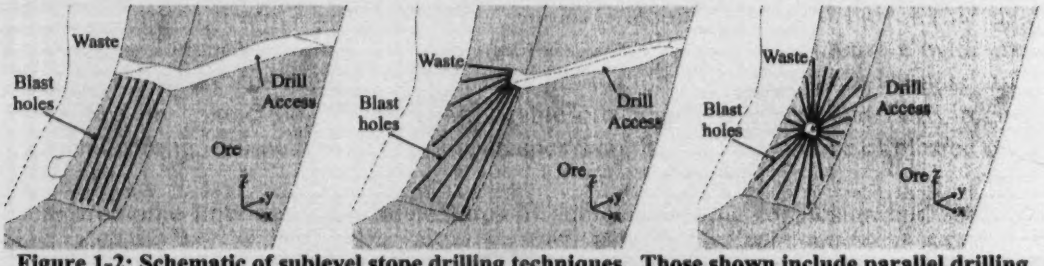

Figure 1-2: Schematic of sublevel stope drilling techniques. Those shown include parallel drilling (left), fan drilling (center) and ring drilling (right).

VCR mining is similar to sublevel stoping; the major difference is in blast design. Stope designs must be very regular with similar upper and lower level geometry. This requirement is due to the blasting procedure used. Parallel or fan drilling is used to drill from the upper level to the lower level. Blast charges are placed at the base of the holes to blast off layers of ore. A percentage of the blasted ore, sometimes equal to its swell factor, is extracted afterwards. This is repeated until the stope has been completely blasted and extracted, see Figure 1-3.

Raise mining can be used as a more selective alternative to sublevel stoping and VCR. This method uses a raise at one end of a stope, a raise-climber and horizontal drilling to blast off slices of ore. The raise climber is essentially an elevator used as a drilling platform. Figure 1-4 shows a long section of raise mining.

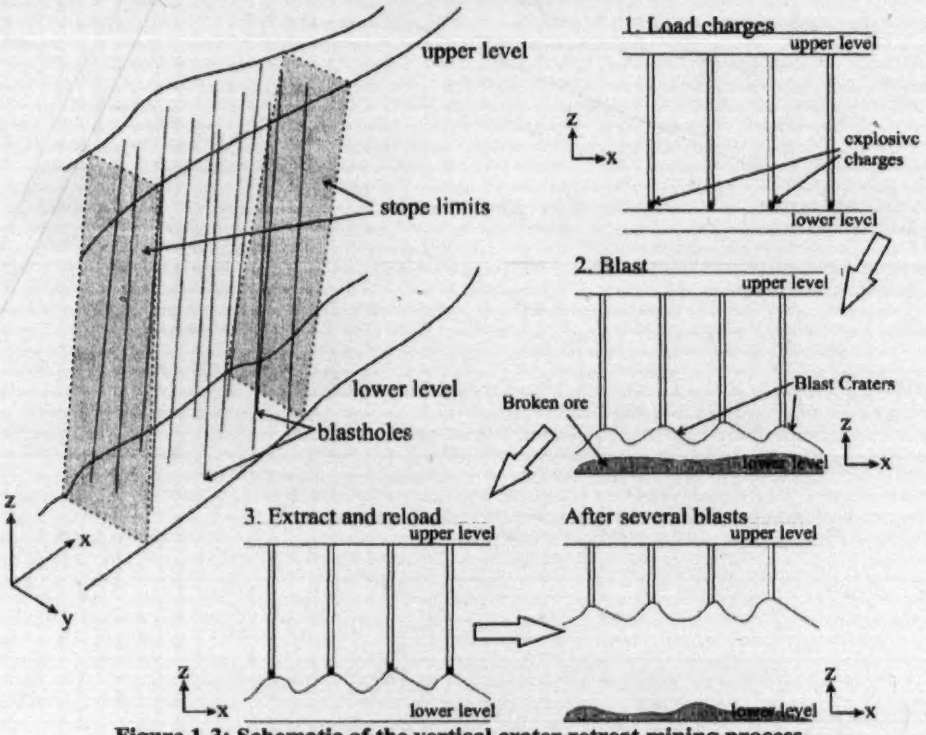

Figure 1-3: Schematic of the vertical crater retreat mining process.

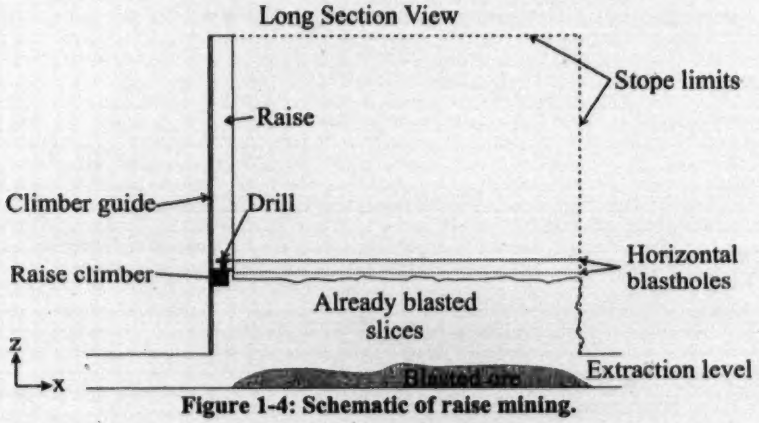

Figure 1-4: Schematic of raise mining.

Cut-and-fill Stoping

Unlike sublevel stoping and its variants, cut-and-fill is a highly selective mining method with lower production rates. The degree of mechanization depends on the size of openings and can range from highly mechanized to the use of hand operated tools. Mining progresses by levels in either an overhand or underhand fashion with each level being backfilled prior to mining the next. Ore extraction is a mechanized process where ore is scooped and transported to vertical or inclined ore-passes leading to a main haulage and extraction level. Applicable ore bodies are steeply dipping vein deposits or those with very irregular geometry. These two deposit types along with weak or unstable rock conditions make cut-and-fill an attractive method.

Many variations of cut-and-fill exist including overhand, underhand, post-pillar, and drift-and-fill stoping [9]. Only drift-and-fill will be discussed here, but Chapter 19.1 in the *SME Mining Engineering Handbook* as well as Chapter 11 in *Introductory Mining Engineering* by Hartman and Mutmansky can be reviewed. Drift-and-fill stoping is typically utilized when mining conditions are particularly poor and the ore body is irregular. Stopes take on the form of drifts and each level consists of several of these. Drifts are mined successively and backfilled with cemented backfill for increased stability. Figure 1-5 shows a schematic of a level plan.

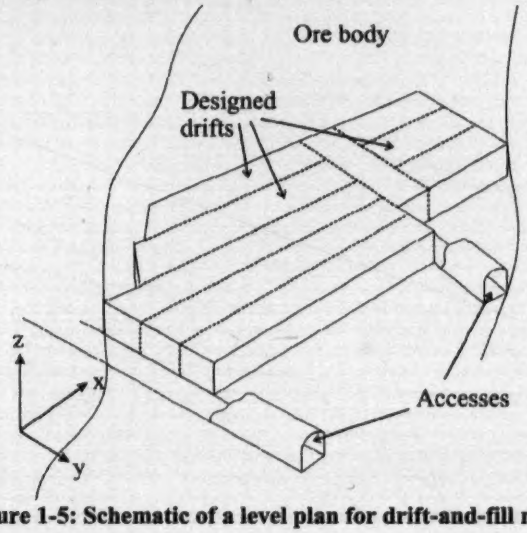

Figure 1-5: Schematic of a level plan for drift-and-fill mining.

Chapter 2 Numerical Optimization

Development of computational techniques for solving various problems is an ever growing field. These problems are typically infeasible for an individual to solve by hand. Creating algorithms to do so, however, is an achievable task with the computing power available today. Numerical optimization algorithms come in many forms and categories. Those of interest are intended to minimize or maximize a function over some space of input parameters. Given a function F(x) for example, the goal is to find its minimum over the parameter space defined by x.

The function in Figure 2-1 has only one parameter and could be minimize rather easily. However, this ease of finding a solution diminishes as the number of parameters increases and the function becomes more complex. In many cases, one cannot even observe what the function looks like. Techniques required to globally optimize a function of many parameters become more creative and interesting, but also more challenging to understand and apply.

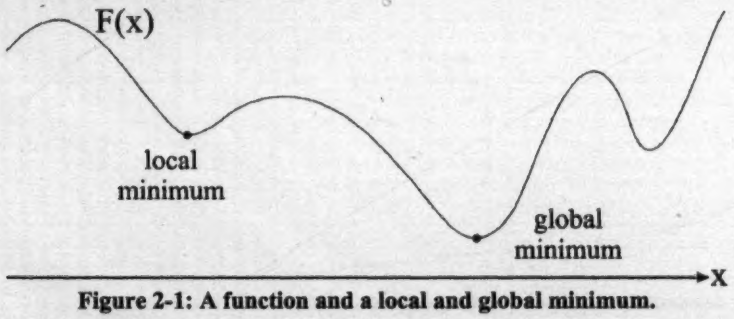

Figure 2-1: A function and a local and global minimum.

Two categories of parameters exist and are applicable to stope geometry and stope sequence optimization. Parameters can be continuous as in Figure 2-1 or discrete [8]. In the later case, problems may be combinatorial and there are a variety of algorithms for solving them. Stope geometry optimization exists in a continuous parameter space in that its location and geometry are defined by Cartesian coordinates. Stope sequencing is a combinatorial optimization problem: each mining event takes on a discrete value describing its order of execution. Only a few of the many available optimization techniques are touched on in this chapter. More information can be obtained from applicable references including [1 to 5, 7, 8, 12, 15, 16 and 17].

2.1 Optimization with Continuous Parameters

A special case of optimization along a continuum is the single parameter or univariate case. Acquiring an understanding of how algorithms work in one dimension supports the development of multidimensional techniques that are impossible to visualize. Some continuous variable optimization methods include line-search, gradient-based methods such as gradient descent and Newton's method, bracketing, and expectation-maximization [1 and 8]. Discussing these techniques is an extensive exercise.

The methods mentioned are applicable primarily for fairly well defined functions with readily accessible gradients. An optimization technique that can deal with an infinite number of options for stope design and geometry as well as for the different mining methods and ground conditions is required. Adaptable optimization methods such as simulated annealing and genetic algorithms are suitable for this problem.

2.1.1 Simulated Annealing

Simulated annealing is based on laws of physics. The underlying process of annealing is intended to make materials such as glass and steel less brittle [15]. These materials are heated to a temperature that is maintained for some time and then cooled slowly and uniformly. The outcome is a more ductile material less susceptible to failure. Relating this information to simulated annealing, the internal structure of the material is analogous to a function to be optimized or made less "brittle". The cooling rate is analogous to the mechanism responsible for accepting or rejecting changes made to the parameter space.

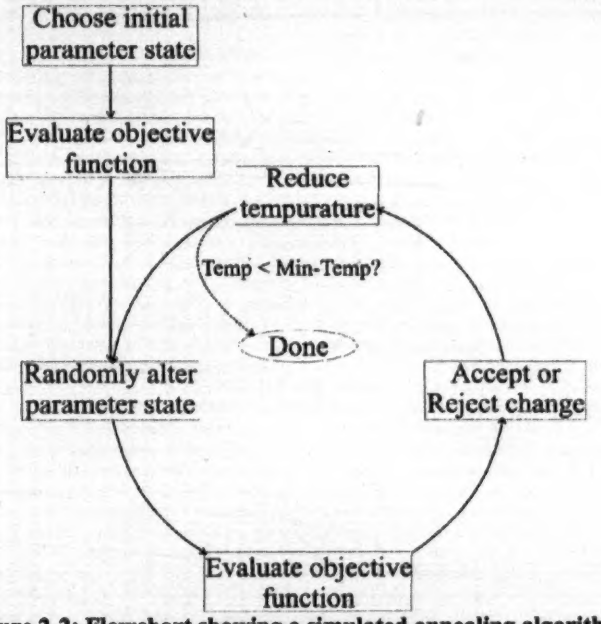

Figure 2-2: Flowchart showing a simulated annealing algorithm.

Simulated annealing involves three main steps: (1) choose initial parameter values; (2) evaluate the objective function for those parameters; (3) randomly alter current parameter state; (4) re-evaluate the objective function and accept or reject the change made in step 3; (5) consider reducing the temperature and return to step 3. Figure 2-2 shows a flowchart of this process. Details of this algorithm applied to stope geometry optimization will be provided in Chapter 3.

Accepting or rejecting changes made to the parameters is done according to a probability. If x_0 is the current parameter state, x_k the new randomly chosen parameter state, F_0 and F_k the objective function evaluated at x_0 and x_k respectively, and T the current temperature parameter, then the probability, *Prob*, can be defined by Equation 2-1. Note that the form of Equation 2-1 is for maximization of a function.

$$Prob = \begin{cases} \exp\left(\frac{F_k - F_0}{T}\right) & \text{if } F_k < F_0 \\ 1 & \text{otherwise} \end{cases} \qquad 2\text{-}1$$

Setting reasonable parameters for simulated annealing is crucial to its success in finding acceptable solutions. Defining a starting temperature and how it is reduced will affect how the acceptance criteria are evaluated throughout the optimization process. Inefficient choices may lead to poor solutions or unreasonably long run-times. Depending on the characteristics of the parameter space, there can be many possibilities in how new parameter sets are chosen. For multidimensional problems, one could change individual parameters, subsets or all parameters simultaneously during annealing iterations. Parameters are also changed by some assigned magnitude. Too small a magnitude can cause long running times or completion prior to reaching a good solution. Too large a magnitude can lead to oscillations in the objective function. Setting up these details of the simulated annealing schedule for stope optimization is discussed in Chapter 3.

2.1.2 Genetic Algorithms

The process of evolution involves entities undergoing structural alteration to acquire better performance in an environment. This is the premise of genetic algorithms (GAs) [5]. A brief description of this type of algorithm is provided below. Algorithm operation is based on mechanisms found in genetics: mutation, crossover and fitness. Other evolutionary features involved are population, chromosomes and genes. A population is a set of solutions to a particular problem. Each solution in the population is a chromosome made up of genes. Genes can be considered characteristics of a chromosome that yield a specific solution. In terms of stope geometry, a chromosome could be a stope design and its genes the parameters used to define the geometry. The stope's value is the solution within a population of other possible designs. Figure 2-3 shows the general algorithm flow.

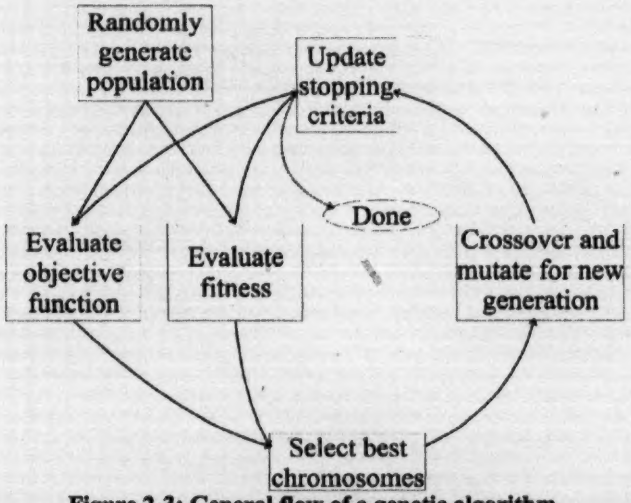

Figure 2-3: General flow of a genetic algorithm.

Like simulated annealing, choosing a starting point amounts to randomly selecting a position in parameter space. For GAs a set of starting points are generated forming the population, each one being a chromosome whose genes are defined by its position in parameter space. The objective function can be evaluated for each chromosome along with its fitness, which is some devised measure of how viable it will be as a parent in the crossover stage. This fitness measure allows parents, perhaps with a poor objective but some good features, to be used in creating the next generation. During crossover and mutation, select genes from each chromosome are used to create a generation of children, which are passed back to the evaluation stage, see Figure 2-4. Defining stopping criteria for GAs can be based on several factors: relative change of the objective function, a desired level of fitness, run-time based, maximum number of iterations, or others.

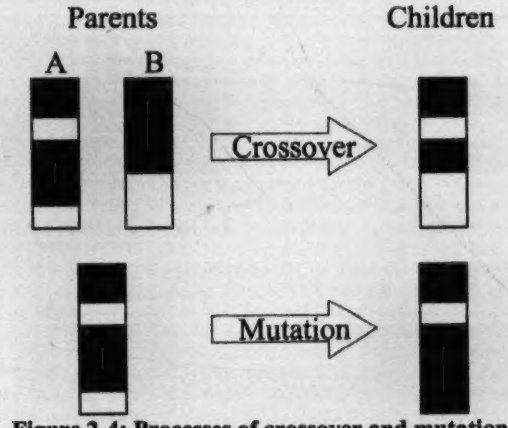

Figure 2-4: Processes of crossover and mutation.

2.2 Optimization with Discrete Parameters

Unlike stope geometry optimization, stope sequencing deals with scheduling; the parameters take on discrete integer values defining their order of execution. This situation can be referred to as a sequential ordering problem. Both simulated annealing and genetic algorithms can be used for these problems, but there are a wide variety of other techniques available. Some of these algorithms include simplex, primal-dual, branch-and-bound, mixed integer programming, dynamic programming, exchanging heuristics, and tabu search. There are many others.

A class of methods of interest for stope sequence optimization is exchange heuristics. Rather than attempting to find new and better sequences from scratch, it may be more efficient to just update an existing sequence by exchanging the order of a pair of stopes. Another method that may prove effective is branch-and-bound. Certain decisions made early on in the sequencing problem can be ignored prior to developing the entire sequence. Ignoring these poor decisions is based on upper and lower bounds of sequence value calculated without having to evaluate an entire sequence.

2.2.1 Exchange Heuristics

Exchange heuristics are a local search technique. Exchanging can be implemented as a 2-exchange, where two variables are altered, or as a k-exchange, where k are altered [5]. Regarding combinatorial optimization, exchanging heuristics essentially involve swapping components of a problem state to yield a better state. Given a current stope sequence, the order in which two different stopes (a 2-exchange) are mined is swapped to provide a sequence with higher value. State-specific information is used by the heuristic to evaluate each possible swap and choose that yielding greatest improvement. One problem with exchange heuristics is that they are local. There is no guarantee of optimality. They are more effective when embedded within another optimization technique such as simulated annealing.

Components of exchange heuristics important to stope sequence optimization are how many stopes in the sequence should be exchanged and how to choose which stopes should be exchanged. Exchanging the sequence of three stopes may yield a better solution that would two for example. Choosing which stopes should be exchanged must be done such that the new sequence is feasible given all constraints of the optimization problem.

2.2.2 Branch-and-Bound

Branch-and-bound utilizes what is called a state-space tree [16]. Each location or node in the tree describes a different problem state and a bound value. Bounds are calculated at each location for determining if a node will lead to a better solution than the current one. If a bound indicates a worse solution, that branch of the tree can be ignored or cut from the traversal process. For stope sequences, the first tree node would contain the null state. The next level of the tree would contain nodes for each stope as if they were chosen first in the sequence; level two would contain nodes for the second stope in the sequence and so on, see Figure 2-5.

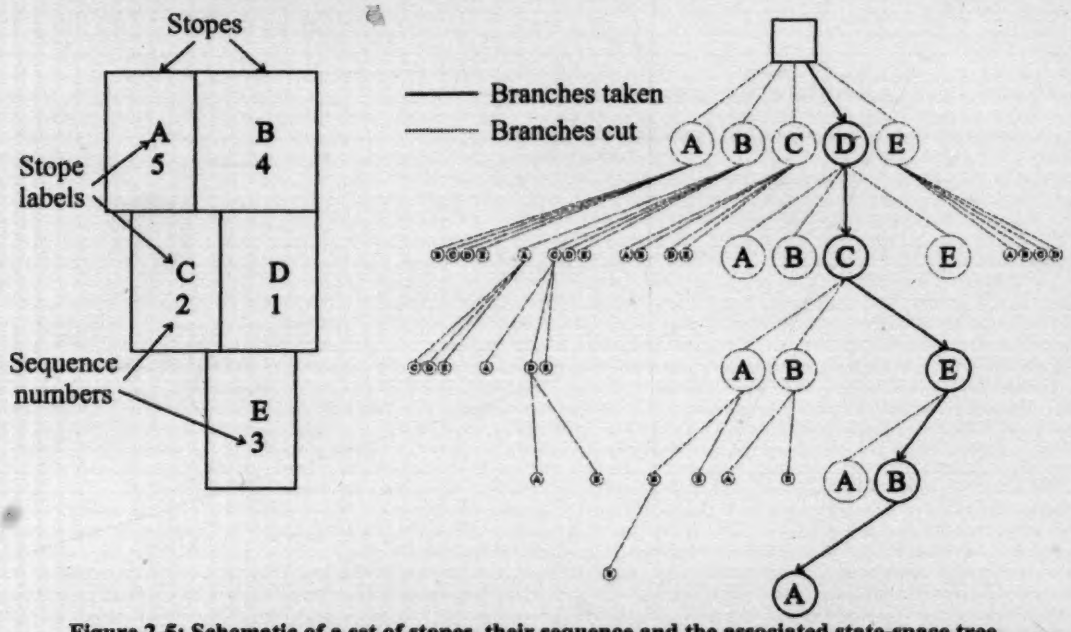

Figure 2-5: Schematic of a set of stopes, their sequence and the associated state-space tree.

For the five-stope example in Figure 2-5 and assuming every possible sequence was feasible, there are 120 total permutations available. One could simply evaluate every branch of the tree and choose the best sequence. When considering ten stopes, however, the total number of permutations is 3,628,800. If a complex objective function is to be evaluated, cutting branches from the tree becomes an important step in branch-and-bound. Panels of stopes could easily exceed 20 stopes (2.43E18 Permutations without constraint) and cutting is critical to finding an optimal solution in a realistic amount of time.

Additional to the requirement of being able to partitioning a solution into exclusive sets, branch-and-bound requires an algorithm for evaluating the upper and lower bounds of the objective function for a given tree node [17].

2.3 Techniques Applied

For generating solutions to stope geometry and stope sequence optimization problems, two numerical optimization techniques were explored. Simulated annealing was applied to both problems because of its ability to handle highly dimensional problems with many local minima. An additional algorithm for stope sequence optimization was applied that chooses sequences logically based on information about the stopes and mining operations. This method avoids the random component of simulated annealing. This logic driven method is similar to branch-and-bound in that stopes are chosen based on them leading to the most optimal sequence. Rather than calculating bounds for each state; however, this algorithm calculates a probability that a choice will provide the optimal solution. These algorithms are described in detail in Chapter 3 for stope geometry optimization and Chapter 4 for stope sequence optimization.

Chapter 3 Stope Geometry Optimization

Improving economic performance of an underground mining operation is an important objective. Modern mines are increasingly marginal; a small percentage increase in profit could mean the difference between moving forward and ceasing operations. One particular area of interest for economic improvement is optimization of stope geometry. The objective is to increase the value from a particular stope by modifying its existing design. Complexity of underlying geology and mineral grades along with their uncertainty make this task difficult to execute by hand. This problem has been implemented computationally to provide a semi-automated solution.

Stope geometry optimization requires several prerequisites: (1) a consistent parameterization of the stope geometry; (2) a pre-specified set of mining constraints; and, (3) a consistent methodology for how changes are made to a stope's existing design. Creating a flexible solution under this framework is non-trivial. A solution is developed for underground mining methods where a stope can be defined as a triangulated solid. This would include methods such as sublevel stoping, vertical crater retreat, raise mining and cut-and-fill (Refer to Section 1.1.3 in Chapter 1). One common factor with all of these methods is their utilization for steeply dipping ore bodies. Shallow dipping or flat lying ore bodies using, for example, room-and-pillar mining would require a different parameterization.

3.1 Parameterizing a Stope

Successfully optimizing a stope to yield improved economics requires an accurate parameterization. A full understanding of location and geometry is important to correctly define constraints. Defining a stope numerically is a computational geometry exercise.

3.1.1 Stope Geometry

For the solution to stope geometry optimization presented here, a stope is defined as a closed triangulated entity, see Figure 3-1. Note that there are several ways in which an ore zone may be identified. In Figure 3-1, it is identified as a single layer sandwiched between two layers of host rock or waste. There are cases where these layers are not well defined and an ore zone may be identified by a cutoff grade.

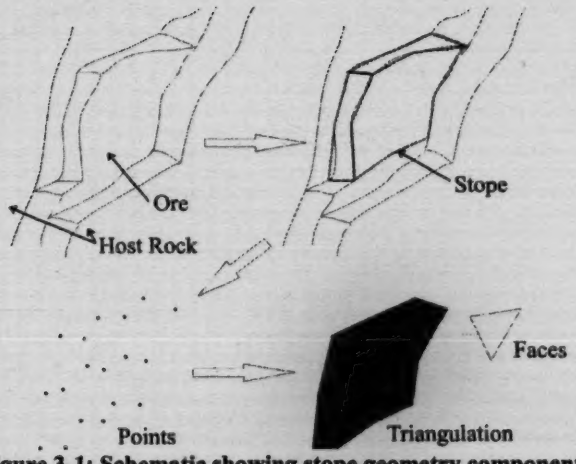

Figure 3-1: Schematic showing stope geometry components.

3.1.2 Constraints

Constraints must be imposed on the methods and procedures used for ore extraction to keep a mine operating safely and effectively. One component of mine design is analysis of the geological, geotechnical and hydrological characteristics. Together, these will help to identify limitations on opening size, blasting vibration and minimal ground support requirements. Here, constraints were implemented in a re-useable fashion. That implies they can be used for numerous mining methods and stope designs. All constraints are geometric and in most cases can be related to features of a particular mining scenario. They include: (1) minimum mining width or thickness, (2) maximum allowable span, (3) minimum and maximum face side-length, (4) maximum allowable deviation across faces, (5) minimum allowable deviation around corners, and, (6) degrees of freedom for individual vertex movement. These are shown in Figure 3-2.

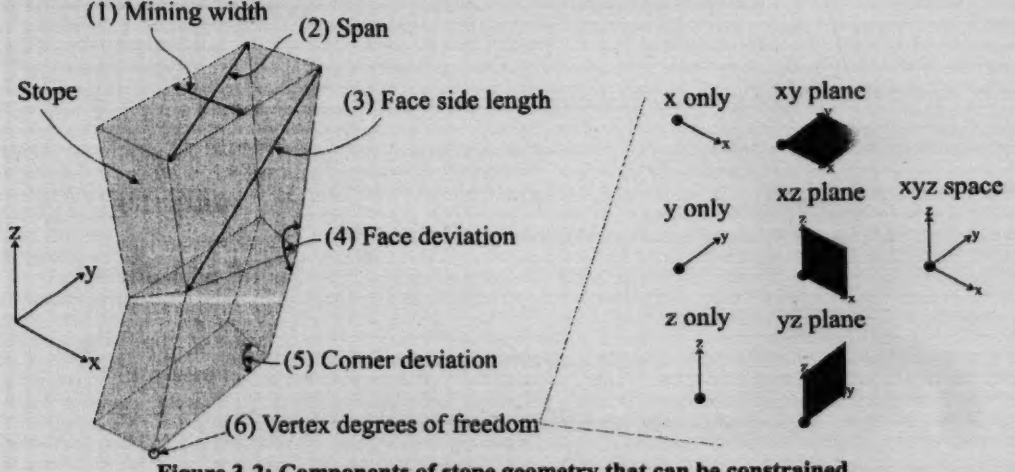

Figure 3-2: Components of stope geometry that can be constrained.

The first constraint, minimum mining width or thickness, was devised based on two possibilities: degree of mechanization and gravity fed ore extraction systems. Depending on the degree of mechanization of an operation, the amount of room needed to operate effectively must be maintained. This is important for methods where development openings, likely made in ore, are of similar dimension as the stope to be drilled. The dimension of interest is often perpendicular to the deposit strike. Another case is cut-and-fill mining where a stope is developed as a drift, see Figure 3-3. All equipment for drilling, blasting and extracting is required to enter the stope. In gravity fed ore extraction systems, a minimum thickness is required to avoid ore hang-ups and ensure smooth flow of ore. Determining a value for minimum thickness might be based on degree of fragmentation of ore from blasting. The largest fragment should flow comfortably through the stope, see Figure 3-4.

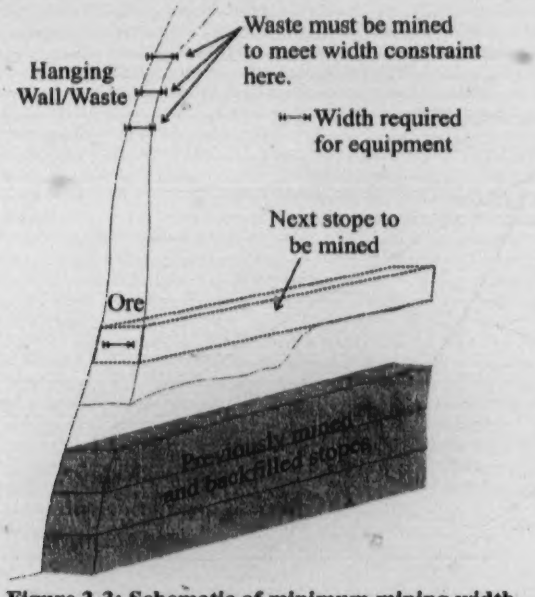

Figure 3-3: Schematic of minimum mining width constraint as applied to cut-and-fill mining.

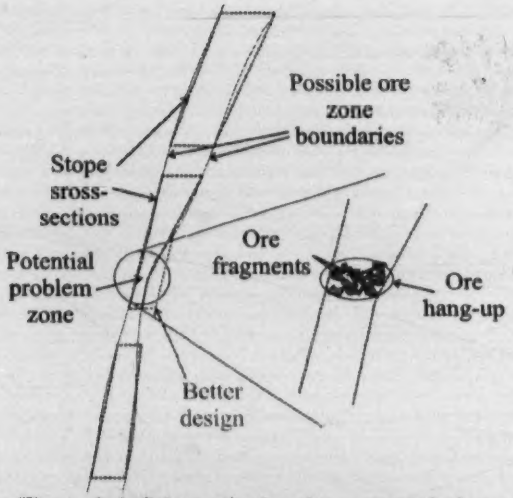

Figure 3-4: Schematic showing a stope design leading to ore hang-up problems.

Maximum allowable span is based on ground stability. From a geotechnical analysis, this is determined as the maximum allowable opening size before collapse. Collapse can range from minor, where rock gradually unravels from the back and walls to an extreme rockburst. As opening size is increased, the weight of the back is transferred to the surrounding walls. Too large an opening will transfer high amounts of stress into the walls leading to failure. Typical design considerations and failure mechanisms are given in Table 3-1 for hard rock mining.

The third constraint is also related to failure mechanisms and stability. Limiting the size of each face describing a triangulated stope can be used to impose maximum overall face dimensions of that stope. One empirical stability analysis method, the Mathews Stability Graph [14], uses hydraulic radius of a face (Equation 3-1) along with features of the rock mass in question to infer if that face is safe to expose.

$$r_H = \frac{w \cdot h}{2(w+h)} \qquad 3\text{-}1$$

Where r_H is the hydraulic radius, w is the width of a face and h is the height. By limiting the size of each triangular face, the hydraulic radius is effectively limited as well. This prevents a stope from growing into an unstable size during optimization, see Figure 3-5. It is also shown that different stope designs can yield different results for the maximum face size dimension constraint. This does not ensure, however, that the same stope design and face size constraint will not lead to different overall face dimensions with optimization. The primary purpose of a minimum constraint on face dimensions is to prevent a face from becoming too small or being pinched out completely from an initial stope design. Depending on how constraints are imposed, it may also aid in maintaining a minimum mining width.

Table 3-1: Potential problems, parameters and analysis for underground hard rock mine openings (paraphrased from [13]).

STRUCTURE	TYPICAL PROBLEMS	CRITICAL PARAMETERS	ANALYSIS METHODS	ACCEPTABILITY CRITERIA
Pillars.	Progressive spalling and slabbing of the rock mass leading to eventual pillar collapse or rockbursting.	• Strength of the rock mass forming the pillars. • Presence of unfavourably oriented structural features. • Pillar geometry, particularly width to height ratio. • Overall mine geometry including extraction ratio.	For horizontally bedded deposits, pillar strength from empirical relationships based upon width to height ratios and average pillar stress based on tributary area calculations are compared to give a factor of safety. For more complex mining geometry, numerical analyses including progressive pillar failure may be required.	Factor of safety for simple pillar layouts in horizontally bedded deposits should exceed 1.6 for "permanent" pillars. In cases where progressive failure of complex pillar layouts is modeled, individual pillar failure can be tolerated provided that they do not initiate "domino" failure of adjacent pillars.
Crown pillars.	Caving surface crown pillars for which the ratio of pillar depth to stope span is inadequate. Rockbursting or gradual spalling of overstressed internal crown pillars.	• Strength of the rock mass forming the pillars. • Depth of weathering and presence of steeply dipping structural features in the case of surface crown pillars. • In situ stress levels and geometry of internal crown pillars.	Rock mass classification and limit equilibrium analyses can give useful guidance on surface crown pillar dimensions for different rock masses. Numerical analyses, including discrete element studies, can give approximate stress levels and indications of zones of potential failure.	Surface crown pillar depth to span ratio should be large enough to ensure very low probability of failure. Internal crown pillars may require extensive support to ensure stability during mining of adjacent stopes. Careful planning of mining sequence may be necessary to avoid high stress levels and rockburst problems.
Cut-and-fill stopes.	Falls of structurally defined wedges and blocks from stope backs and hanging walls. Stress induces failures and rockbursting in high stress environments.	• Orientation, inclination and shear strength of structural features in the rock mass. • In situ stresses in the rock mass. • Shape and orientation of stope. • Quality, placement and drainage of fill.	Numerical analyses of stresses and displacements for each excavation stage will give some indication of potential problems. Some of the more sophisticated numerical models will permit inclusion of the support provided by fill or the reinforcement of the rock by means of grouted cables.	Local instability should be controlled by the installation of rockbolts or grouted cables to improve safety and to minimize dilution. Overall stability is controlled by the geometry and excavation sequence of the stopes and the quality and sequence of filling. Acceptable mining conditions are achieved when all the ore is recovered safely.
Non-entry stopes.	Ore dilution resulting from rockfalls from stope back and walls. Rockbursting or progressive failure induced by high stresses in pillars between stopes.	• Quality and strength of the rock. • In situ and induced stresses in the rock surrounding the excavations. • Quality of drilling and blasting in excavation of the stope.	Some empirical rules, based on rock mass classification, are available for estimating safe slope dimensions. Numerical analyses of stope layout and mining sequence, using three-dimensional analyses for complex orebody shapes, will provide indications of potential problems and estimates of support requirements.	A design of this type can be considered acceptable when safe and low cost recovery of a large proportion of the orebody has been achieved. Rockfalls in shafts and haulages are an unacceptably safety hazard and pattern support may be required. In high stress environments, local distressing may be used to reduce rockbursting.
Drawpoints and orepasses.	Local rock mass failure resulting from abrasion and wear of poorly supported drawpoints or orepasses. In extreme cases this may lead to loss of stopes or orepasses.	• Quality and strength of rock. • In situ and induces stresses and stress change in the rock surrounding the excavations. • Selection and installation sequence of support.	Limit equilibrium or numerical analyses are not particularly useful since the processes of wear and abrasion are not included in these models. Empirical designs based upon previous experiments or trial and error methods are generally used.	The shape of the opening should be maintained for the design life of the drawpoint or orepass. Loss of control can result in serious dilution of the ore or abandonment of the excavation. Wear resistant flexible reinforcement such as grouted cables, installed during excavation of the opening, may be successful in controlling instability.

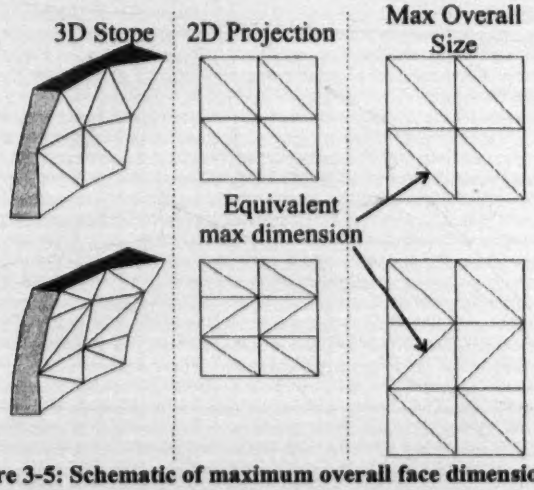

Figure 3-5: Schematic of maximum overall face dimensions given constrained triangulated face dimensions.

Constraints on the maximum deviation across faces and minimum deviation around corners were devised based on capability of mining equipment and blasting. When considering drilling equipment used to drill blast holes, high deviations across a face may be impractical, see Figure 3-6. High deviations may also lead to ore hang-ups as explained earlier. For the case shown in Figure 3-6 where drilling takes place from only the top or bottom, a simpler stope design would be desirable and such high deviations across the faces would not be permitted. Deviation around corners maintains the same principles in an aerial plane. Designs that cannot be drilled or blasted are not considered. Tight corners on the perimeter of a stope are likely infeasible in terms of drilling and blasting, see Figure 3-7.

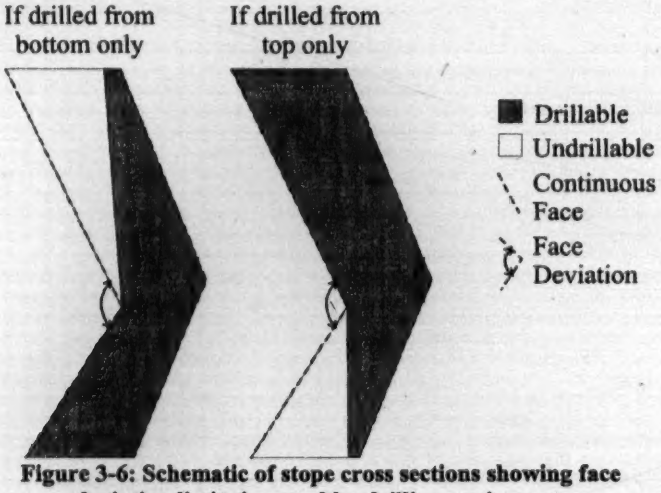

Figure 3-6: Schematic of stope cross sections showing face deviation limits imposed by drilling equipment.

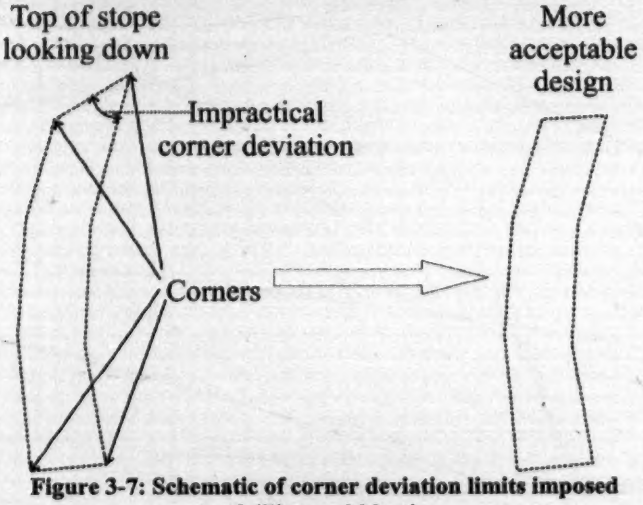

Figure 3-7: Schematic of corner deviation limits imposed by drilling and blasting.

The final constraints are the degrees of freedom permitted for individual vertices defining a stope. Each vertex can be given different movement freedom. It is common for mine operations to have development and extraction levels at fixed depths. Stope designs in this case will have a top and base set of vertices that are restricted to the horizontal plane when being perturbed. Another possible use of vertex movement constraints is neighboring stopes or the requirement for pillars between stopes. This may restrict vertices to movement in vertical cross-section planes of the deposit. Vertices may be completely locked if necessary.

3.2 Block Model Clipping

To evaluate a stope's value, it must be intersected with an economic block model. Other costs such as those related to ground support can be evaluated based on the stope geometry and the surrounding geological and geotechnical features. Clipping is very common practice in computer graphics applications where objects and information are clipped against a view frustum [11]. Many clipping algorithms are made available by computer graphics and computational geometry research. One of the simplest to implement is the depth-buffer method. Relating stope geometry optimization to computer graphics, any space inside a stope object is considered visible and the stope is equivalent to a view frustum. This comparison is easy to understand in two dimensions, see Figure 3-8.

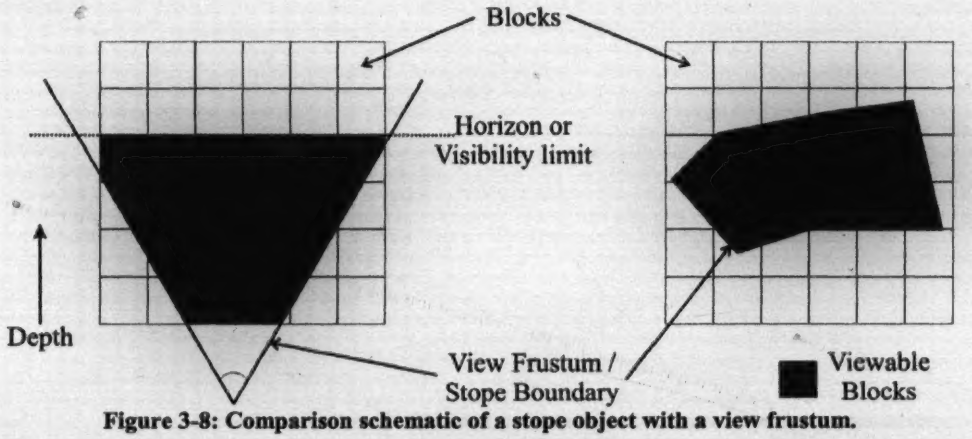

Figure 3-8: Comparison schematic of a stope object with a view frustum.

An additional aspect of block model clipping that is shown by Figure 3-8 is dealing with blocks that are not fully inside the stope. In determining the total value of a stope, only a fraction of the value of these partial blocks should contribute. Three stages of clipping take place to evaluate a stope's value: primary using bounding boxes, secondary using a z-buffer type algorithm, and tertiary to evaluate partial blocks.

3.2.1 Primary Clipping

Primary clipping involves creating an axis-aligned bounding box around the stope object. In this case, the axes would also align with the block model. The entire block model is compared against the bounding box and only those inside are considered for secondary clipping. This is a very fast operation and can be accomplished with grid-block index calculations. For a regular grid (Figure 3-9) defined by an origin as the minimum of all three coordinate axes (x_0,y_0,z_0), cell dimensions, *dx*, *dy* and *dz*, and number of cells along each axis, *nx*, *ny* and *nz*, calculating indices from point locations can be accomplished by Equation 3-2.

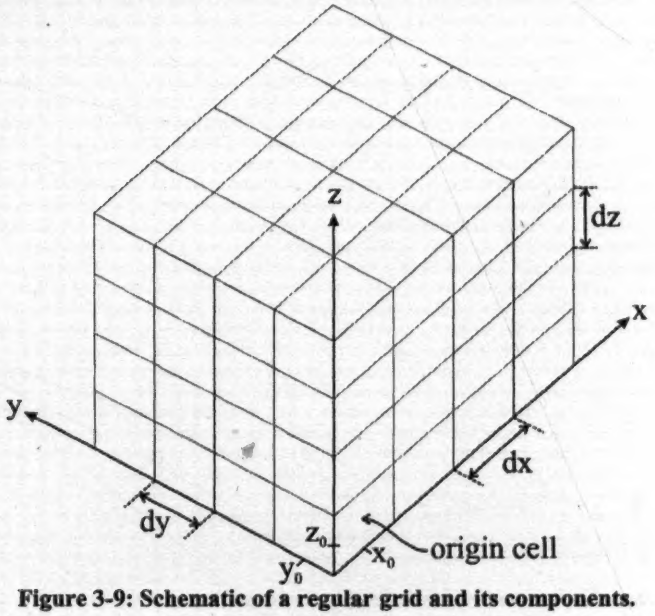

Figure 3-9: Schematic of a regular grid and its components.

$$x_{index} = \frac{x - x_0}{dx} + 1$$

$$y_{index} = \frac{y - y_0}{dy} + 1 \qquad 3\text{-}2$$

$$z_{index} = \frac{z - z_0}{dz} + 1$$

Calculating the indices of blocks coinciding with the minimum and maximum corner-points of the stope bounding box immediately provides the subset of blocks for secondary clipping.

3.2.2 Secondary Clipping

The second stage of block model clipping is to determine which blocks of the subset past from primary clipping are at least partially inside the stope object. A z- or depth-buffer is sometimes used in computer graphics for visible surface detection. Depth values, where depth is measured along line of sight, are calculated over polygons and surfaces at individual pixel positions. Pixels are assigned colors based on the surface having the smallest depth value. In terms of defining blocks as in or out of a stope, the line of sight is a vertical line from the top of the stope descending through a block of interest. If the top of the block comes before the base of the stope, it is inside, see Figure 3-10. Note that the base of the stope can be any part that is intersected by the line of sight, not necessarily the actual base.

After secondary clipping, blocks will be given a status as fully inside, partially inside or fully outside. Only those blocks identified as partially inside are considered for tertiary clipping.

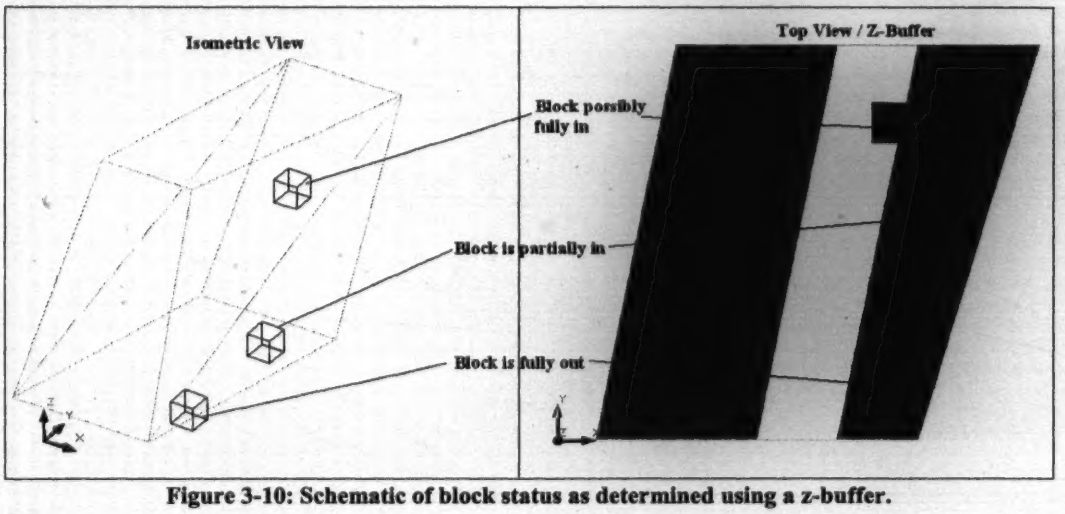

Figure 3-10: Schematic of block status as determined using a z-buffer.

3.2.3 Tertiary Clipping

Determining the fraction of a block inside a stope is important to calculate the stope's value. The actual piece of the block inside the stope is not required. The method used is similar to quadrature for integration over a specified interval. A difference, however, is that refinement is not carried out until the error is acceptably low; it continues a specified amount of times indicated by the user. Refinement of a block is the process of splitting it into eight smaller blocks. For n refinement steps of a particular block, the process of calculating fraction inside is done as follows (refer to Figure 3-11).

1. Refine the block into 8 sub-blocks
2. Define those eight blocks as fully inside, partially inside or fully outside
    a. Add blocks with fully inside status to fraction-in value
3. Pass each block with partially inside status to step 1
4. After final iteration, add 50% of blocks with partially inside status to fraction-in

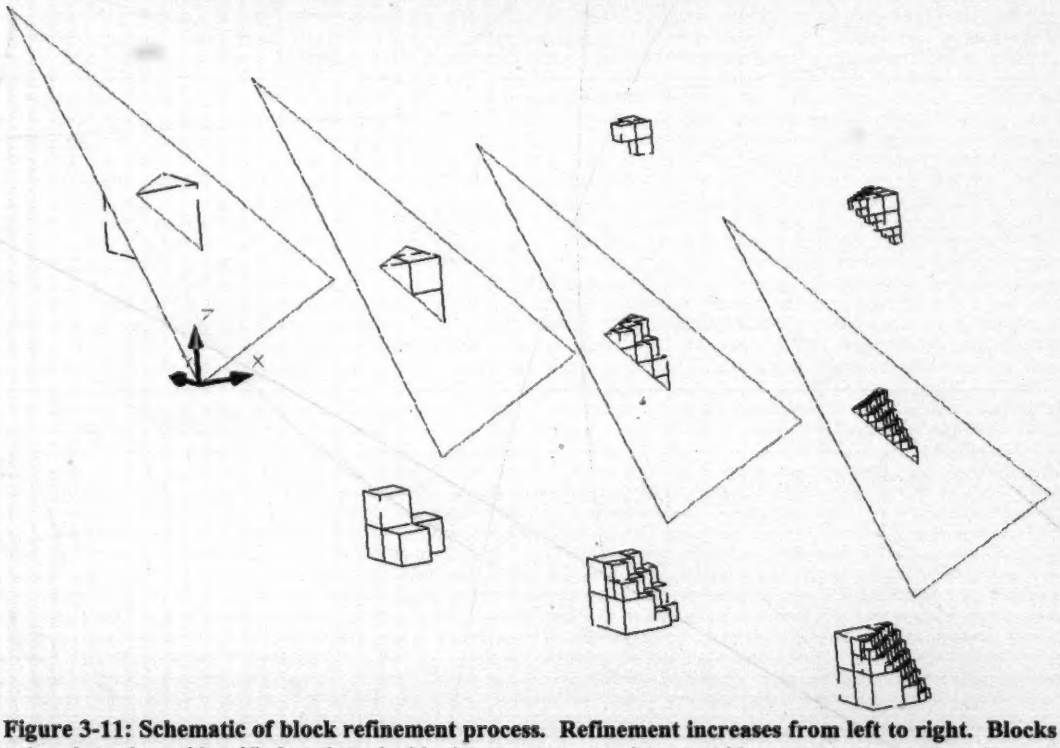

Figure 3-11: Schematic of block refinement process. Refinement increases from left to right. Blocks below have been identified as fully inside, blocks above as fully outside, and blocks on the plane as partial.

3.3 Stope Modification

A stope's design must be changed to see if its value can be improved. One way this can be done is by shifting a vertex making up the design. Ideally, we would like to move all vertices so the stope's value is maximized and all constraints are satisfied. Determining these locations is the goal of optimization. When a vertex is moved, the following operations must be carried out: (1) update the stope object, which includes updating face and vertex information, (2) check all constraints to make sure they have not been violated, (3) clip the block model using the updated stope object, and, (4) evaluate the new stope's economic value.

3.4 Optimization with Simulated Annealing

Maximizing economic value of a stope by modifying its existing design is a high dimensional constrained optimization problem. An infinite number of stopes are possible and the underlying geologic and economic block models can be very complex. Models and parameters also contain uncertainty. Simulated annealing was chosen as an initial optimization technique because it can deal with highly dimensional problems that do not present a clear or even conceivable solution. Recall the steps to simulated annealing outlined in Chapter 2: (1) choose initial parameter values, (2) evaluate the objective function for those parameters, (3) randomly alter current parameter state, (4) re-evaluate the objective function and accept or reject the change made in step 3, (5) reduce the temperature and return to step 3. Iteration is repeated until the temperature is low enough or convergence on a solution that cannot be improved has been made. This process as applied to stope geometry optimization is shown by Figure 3-12.

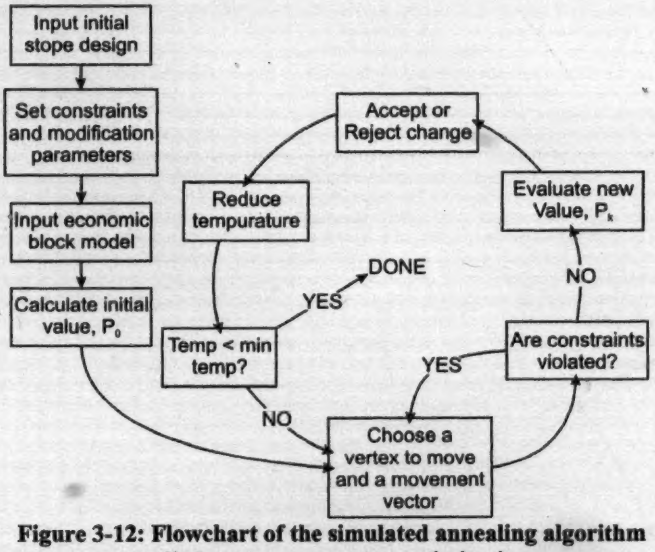

Figure 3-12: Flowchart of the simulated annealing algorithm applied to stope geometry optimization.

Selection of initial parameters involves the initial stope design that is passed to the algorithm as well as any pre-defined constraints. Constraints depend on mining method and operational procedures, geology, and geotechnical analysis results. Evaluating the objective function, which is defined by Equation 3-3, involves clipping the economic model and summing the value of all blocks within.

$$V_S = \sum_{\forall b \in S} f_b \cdot v_b \qquad 3\text{-}3$$

Where V_S is the value of stope **S**, **b** is the set of blocks, f_b is the fraction of a block inside stope **S** and v_b is the value of that block.

Modifying the stope design involves two random processes: choosing a vertex to move and choosing a vector defining where to move it. Both of these are done with a uniform random number generator where numbers are between zero and one. For a vertex, this number is multiplied by the number of vertices plus one and truncated to an integer. A direction vector is chosen as a normally distributed parameter with user-specified starting mean and standard deviation values. Using normally distributed motion vectors was based on there being a higher probability of choosing vectors with smaller length and lower probability of choosing vectors with large magnitudes. Choosing more vectors that make small changes to the stope will allow convergence to local maxima. Occasional drawing of large vectors will give opportunity to step away from local maxima, possibly towards a better solution.

Deciding on a schedule for the annealing temperature can be difficult. Poor selection could result in premature stoppage of optimization or extremely long run-times. For the algorithm developed here, the initial temperature is set as the initial stope value. Since the temperature represents a dollar value, the user can select the algorithm's stopping temperature keeping these units in mind. Improvements in a stope's value of only a few dollars may not be a concern and excess time will not have to be spent searching for a better solution. For every proposed modification that a stope receives, the temperature is reduced. If a modification is rejected, the temperature is reduced slower than if it is accepted. This allows for the algorithm to propose more modifications to move beyond local maxima.

The algorithm developed here was applied to three stope designs. A very simple case was used so that the objective surface could be visualized. A stope from sublevel stope mining and another from cut-and-fill mining were explored as well. More detailed descriptions of the stopes, mining scenarios, and economic models are given in Chapter 5.

Chapter 4 A Framework for Stope Sequence Optimization

Having an optimal stope design in terms of geometry is not the only way to maximize economic value. The operating schedule can be altered. Changing the sequence of operations can be done with various goals in mind including: avoidance of hazardous equipment interactions, minimizing operating costs, blending ore and sending optimal grade to the processing plant, and maximizing potential economic return. Stope sequence optimization maximizes the net present value (NPV) of a set of stopes, which may comprise an entire mine or a smaller portion of it. This problem can be considered a sequential ordering problem: the order in which a set if stopes are mined is reconfigured to maximize NPV. Discounting is accounted for over the mine-life.

The large number of possible scenarios and variables make development of a flexible algorithm difficult. Solving one particular aspect of underground mine schedule optimization is shown below with the intention of creating a framework for formulating solutions to other similar problems. The mining scenario considered involves steeply dipping ore bodies that can be segmented into sets of stopes called panels and each stope having similar properties and constraints, see Figure 4-1. Operations used to access, prepare and mine each stope are considered constant for a given panel.

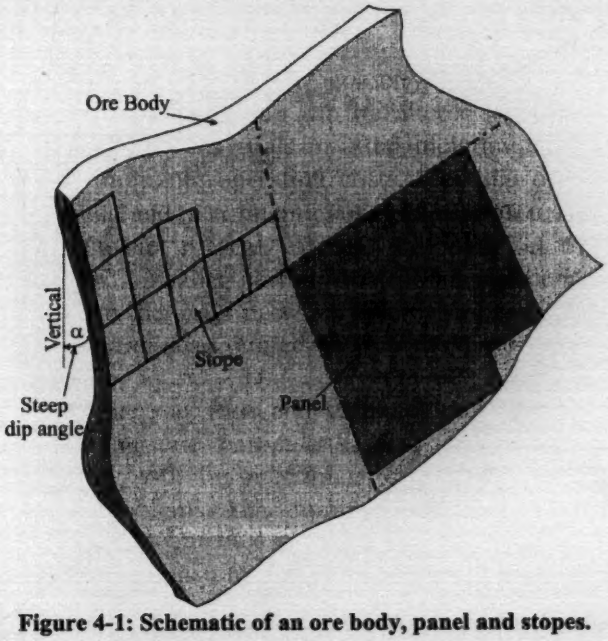

Figure 4-1: Schematic of an ore body, panel and stopes.

A variety of combinatorial optimization techniques as outlined in Chapter 2 exist for the sequential ordering problem. Two methods were explored: (1) simulated annealing as a stochastic method; and (2) a logic-driven algorithm that will be referred to as

probabilistic decision making (PDM). Simulated annealing was chosen because of its relative ease of implementation and to provide a benchmark for other attempted algorithms. The second method used logic to determine a sequence offering optimal NPV where choices are made based on probabilities. Before implementing any algorithm, however, several topics must be clarified including how a sequence will be parameterized and represented numerically as well as the calculations required to determine feasibility and evaluate the objective function. Determining feasibility involves scheduling a sequence that can be mined without violating any constraints or mining equipment limitations.

4.1 Parameterizing a Sequence

Adequately setting up all variables, characteristics and constraints of a sequence of stopes involves parameterizing its sub-components. These include the panel of individual stopes, development openings, mining equipment, crews and procedures. Equipment, crews and procedures can be described as mining operations. Components are used together to characterize a mining schedule or sequence of operations for extraction of a complete panel.

4.1.1 Stope Parameterization

Characterizing a stope for sequencing is similar to that for stope geometry optimization; however we do not intend to change the stope design. There are no constraints associated with stopes on an individual basis. Constraints are integrated with operations. The value of each stope must be known. Other parameters describing stope location and size, neighboring stopes, and other nearby openings and geometry-dependent costs such as rock support are required.

To calculate the NPV of a stope, we need to know the following information: its volume or tonnage and undiscounted economic value, any fixed and pre-mining costs, stope preparation time, the time mining starts on that stope and how long it takes to mine, the cost associated with mining the stope, time and cost to add backfill, and how long it must settle after backfilling. Some of these parameters are calculated after a stopes position in a sequence is known as they are dependent on operation parameters. An assumption is made that ore is processed immediately after mining.

Most stope parameters can either be directly input or calculated prior to optimization. Volumes, tonnages and undiscounted economic values may be acquired from the stope geometry optimization as described in Chapter 3. Fixed and pre-mining costs are associated with preparing the stope for mining. These might include costs for developing drilling platforms and extraction systems, installing rock support, drilling, loading explosives and blasting. Preparation time and settling time after backfilling can be pre-determined as they will likely depend on volume or tonnage and stope geometry, all of which remain constant.

Some parameters must be calculated as optimization progresses. The time mining begins on a stope depends on its sequence order. Time and cost for mining will depend on equipment operating parameters, extraction methods, available routes for haulage and locations for dumping at the time of mining, and possibly the occurrence of equipment interactions. Backfilling time and cost may depend on similar operating characteristics. Transportation of backfill to a mined stope will depend on available routes and equipment characteristics.

4.1.2 Development Openings

Mine development openings that are utilized for ore extraction, haulage, and equipment transport include drifts, ramps and crosscuts. Each of these components should be fully characterized by an existing mine design. Cross-sectional geometry of each of these openings is fixed. Ramps will involve an additional fixed parameter, that being a grade or maximum grade for equipment traversal. Other features that may be important that are not fixed include: rock-type mined along the openings and ore grade mined along the opening if development is in ore. This information can be determined by merging the mine design with geology and grade block models. Rock-type and its quality (perhaps designated using RMR) are important as they affect drilling equipment, blast design, explosives used and rock support.

4.1.3 Parameterizing Operations

Characteristics of individual pieces of equipment to be used for various tasks are required for calculating operating costs and time. Tasks can be divided amongst different crews, where a crew constitutes a set of equipment and manpower for a specific task. For most mining methods, crews can be segregated into those for main development, those for production development, and those for ore or waste extraction and haulage. In some cases, ore extraction and haulage may be divided into two separate crews. Depending on the dimensions of openings, which will reflect degree of mechanization of a mine, equipment can range from handheld to large remote controlled vehicles. Using drilling equipment as an example, mines with small unstable openings likely use jacklegs whereas those with larger openings and higher production demands might use a drill jumbo, see Figure 4-2.

Main development crews would consist of equipment for rock support, development drilling and blasting along drifts, ramps and crosscuts. Blasted rock must also be extracted from these openings as progress in made. This can be accomplished with a waste extraction and haulage crew. A basic procedure for development (Figure 4-3) involves the following steps: (1) drill off a set of blastholes, constituting a round, (2) load explosives and blast that round, (3) extract and haul blasted material, and (4) install any required rock support. Production crew operations include: (1) mine drill and extraction levels for the stope, (2) install any required rock support, (3) drill production holes, and (4) load explosives and blast the stope.

In terms of stope sequencing, the time to drill off development rounds and stopes, install rock support, and load and blast explosives are needed. Costs of performing these operations are required. In assessing how long it will take to advance one round in a development opening, the following parameters are involved: number of holes to drill and their length, summed up as total length to drill, drill rate per unit length, explosive loading and tying-in time, blasting time, time to extract blasted rock, which will be parameterized below, and time to install rock support. The same set of parameters is required for preparing a stope for extraction. Determining the total length to be drilled comes from the mine design. Drill rate may be calculated or acquired from equipment information sheets and drilling conditions. Tying-in a set of blast holes refers to adding blasting caps, boosters and fuses and connecting this to a detonation mechanism. Installing rock support may involve setting up wire mesh, installing rock bolts and/or split sets and spraying shotcrete. Each of these processes takes time and will depend on the surface area to be covered and the integrity of the surrounding rock.

Figure 4-2: Jackleg drill (top) (source: www.mg.mtu.edu accessed July 20, 2006), and drill jumbo (source: www.atlascopco.com accessed July 20, 2006).

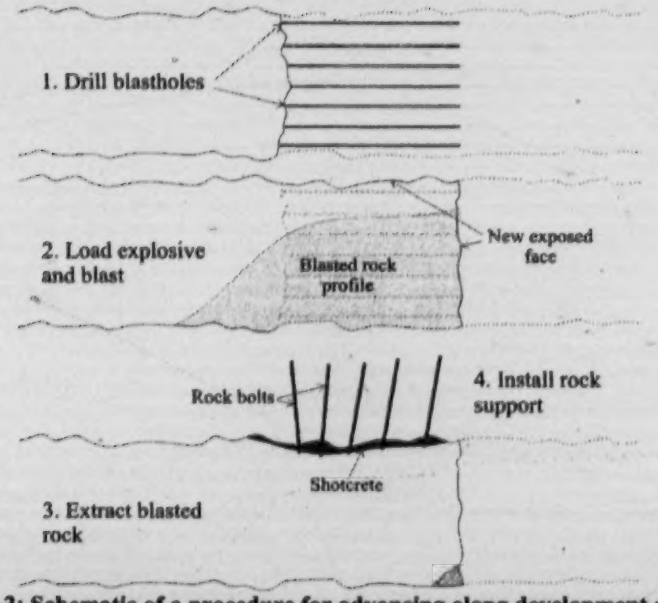

Figure 4-3: Schematic of a procedure for advancing along development openings.

Ore extraction and haulage crews consist of equipment for scooping ore from a stope and transporting it to a dump location. In some cases, these two operations are accomplished with the same piece of equipment called a load-haul-dump (LHD) or scoop-tram, see Figure 4-4. Where haulage distances are substantial, mine trucks may be used. Again, time and cost are important to stope sequencing. To evaluate time, a scoop tram's capacity is required along with how long it takes to acquire a load, transport it to a dump location and dump it. Transporting the load will require information about equipment acceleration rate, straight line traverse speed, cornering speed and braking rate along with the path leading from the stope to the dump location. If ore is transferred to a mine truck, this information is needed for that equipment as well.

Figure 4-4: A scoop tram (top) and mine truck (bottom) (source: www.atlascopco.com accessed July 20, 2006).

Evaluating the cost of operations requires cycle time, the equipment used, and cost per hour. Equipment cost per hour depends on many factors including fuel usage, disposable component replacement (tires, drill bits, etc), maintenance costs and manpower costs. This is not the focus of this thesis. All information regarding equipment and crews is used to decide what can be accomplished at a particular time.

4.1.4. Constraints

Feasibility is important in scheduling mine operations. To ensure safe operation throughout a panel of stopes, various constraints will be imposed affecting which mining tasks can take place, where they can occur and at what time. The basic procedure to extract a set of stopes can be summarized as follows:

1. Make a stope accessible by completing any main development leading to it. Any required haulage drifts must also be mined.
2. Prepare the stope for mining. This includes mining drill and extraction levels, developing the extraction system, and drilling and blasting the ore.
3. Extract all ore from the stope and haul it to a dump location.
4. Backfill the stope.
5. Repeat the procedure for all remaining stopes.

This procedure may be executed on several stopes simultaneously depending on the number of crews available for each operation. For one development and one extraction crew, the above procedure is shown by Figure 4-5. Many possible constraints exist. They depend on the mining environment. The following is a list of potential constraints:

- Considering a stope being mined:
  - Once a stope is completely mined, it may stand open for some maximum time before requiring backfill.
  - Neighboring stopes must stand as pillars until the current stope is backfilled and settled. The arrangement of which stopes act as pillars may change depending on the arrangement of those standing open.
  - Simultaneous mining (development and/or production blasting) cannot take place near the current stope on the same level and/or adjacent levels.
- Regarding timing of events:
  - No significant gaps in ore extraction can take place (i.e. at least one stope must be undergoing extraction at all times).
  - Wait time between ore transfers (LHD to mine truck) must be kept minimal.
  - Equipment must complete an assigned task before moving on to another. For example, a development crew cannot partially complete one round and then transfer to a different drift heading.

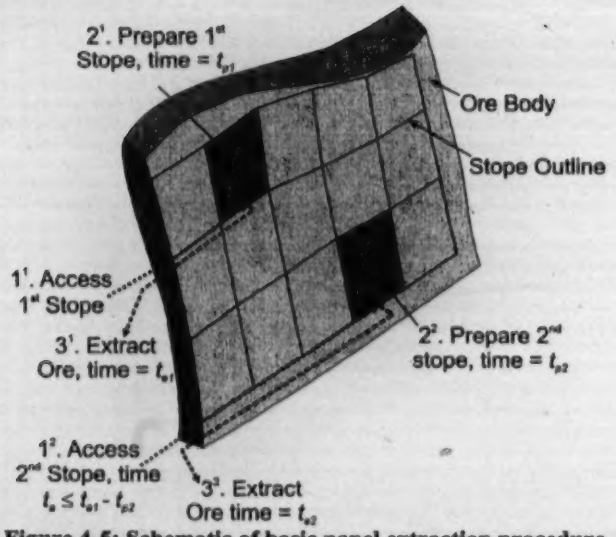
2¹. Prepare 1st Stope, time = t_{p1}
Ore Body
Stope Outline
1¹. Access 1st Stope
3¹. Extract Ore, time = t_{e1}
2². Prepare 2nd stope, time = t_{p2}
1². Access 2nd Stope, time $t_a \leq t_{e1} - t_{p2}$
3². Extract Ore time = t_{e2}

Figure 4-5: Schematic of basic panel extraction procedure.

Clearly defining these constraints will determine what a feasible sequence is. This is very important to the optimization process. Defining too few constraints may cause optimization to choose a sequence that cannot be feasibly mined by operations. Defining too many may prevent an optimal solution from being found.

4.2 Choosing a Feasible Sequence

Choosing a sequence that can actually be accomplished depends on what the equipment can achieve and the constraints imposed. Violating any constraint makes a sequence infeasible. Consider the scenario shown in Figure 4-5 and the following constraints: (1) there can be no breaks in mining of stopes; and, (2) adjacent stopes must remain as pillars. Figure 4-6 indicates which stopes cannot be mined given that the first has already been chosen. If stopes are indexed by level then stope, Figure 4-6 shows that by mining stope 12, stopes 11, 13 and 22 cannot be mined as they must stand as pillars until stope 12 is backfilled and has settled. Stopes 15, 25 and 35 cannot be mined since they cannot be accessed and prepared in the time it takes to mine stope 12. A break in mining would take place otherwise. Regardless of the constraints, this type of logical exercise can be carried out to determine what is feasible given the current state of a sequence.

An advantage of having constraints is they place limits on the number of feasible sequence permutations for a set of stopes. For the scenario in Figure 4-6, having two constraints has limited the number of options from 14 stopes in the unconstrained case to 8 for that particular sequence state. Constraints may also cause entire portions of sequences to be infeasible. Again using the scenario from Figure 4-6 and assuming the current state of the sequence is as shown in Figure 4-7, it can be shown that some sequences lead to dead ends. A state-space tree shows sequence steps that end with one or more stopes being unmineable without violating constraints, see Figure 4-8.

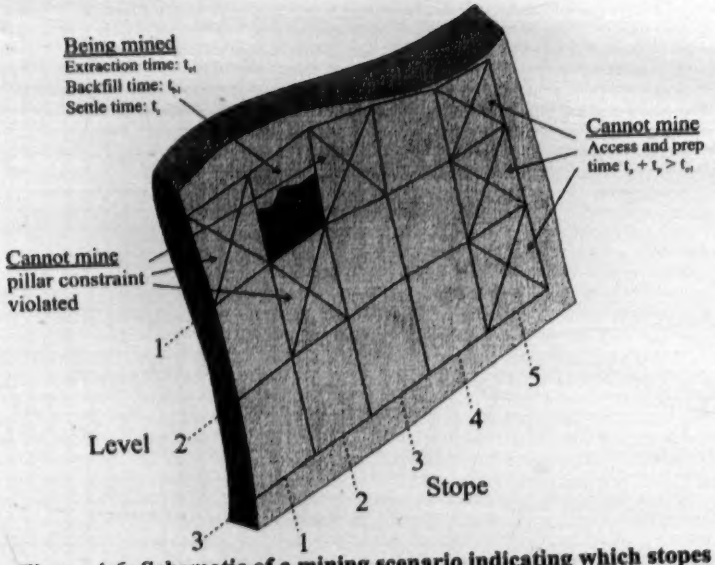

Figure 4-6: Schematic of a mining scenario indicating which stopes are infeasible given constraints.

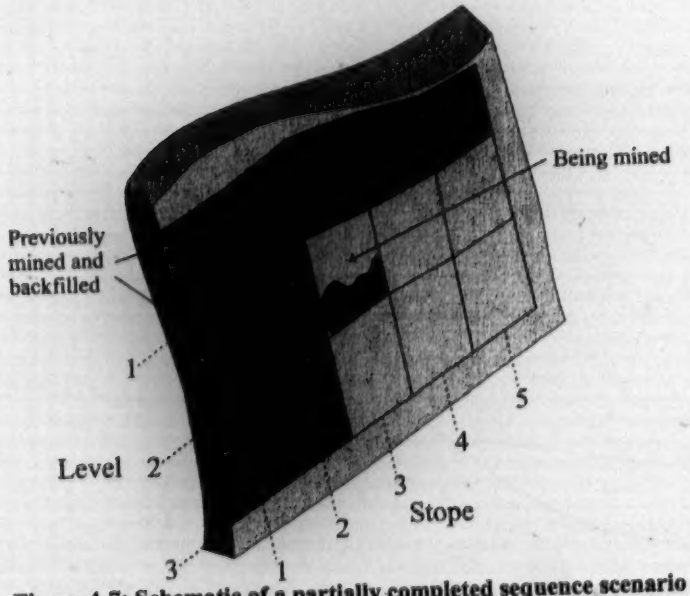

Figure 4-7: Schematic of a partially completed sequence scenario for state-space tree analysis.

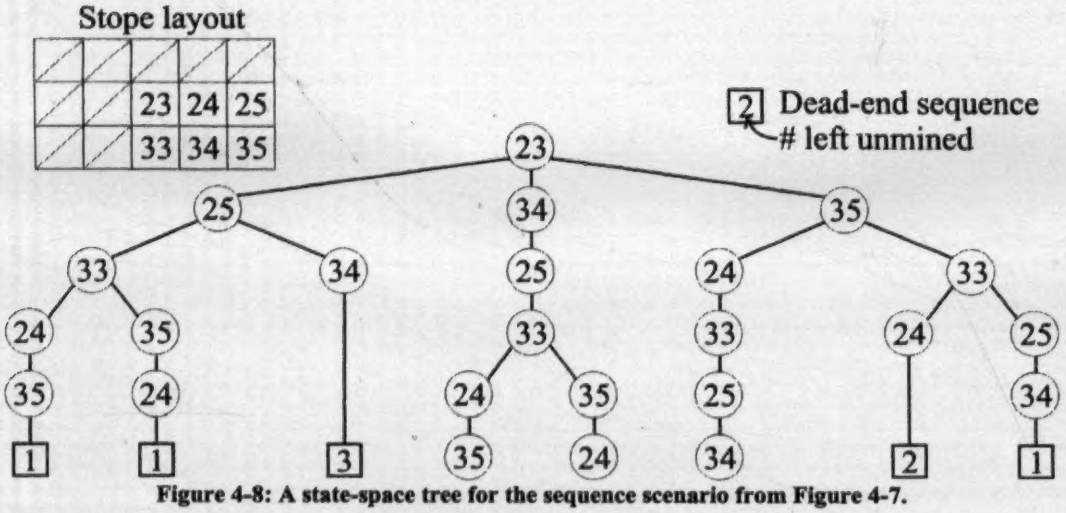

Figure 4-8: A state-space tree for the sequence scenario from Figure 4-7.

4.3 Stope Sequence Scenario

A simple stope sequencing case was used for developing optimization algorithms. This case involves a panel of stopes that can be considered two dimensional. Only one development crew, one stope preparation crew and one extraction crew were considered. Development openings are driven into the panel along strike and below stopes. These openings are also used as haulage drifts leading to a single ore pass. Extracted material is hauled to the ore pass where it is dumped and flows to some secondary extraction system for bringing material to the surface. Figure 4-9 shows an example for this sequence case.

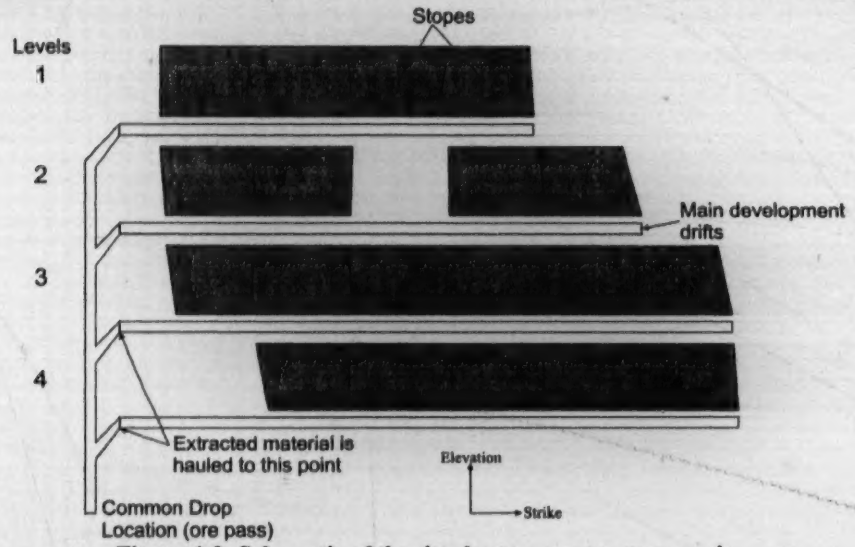

Figure 4-9: Schematic of the simple stope sequence scenario.

Three sequence constraints were applied in this case: (1) once a stope is chosen to be mined, it must be fully prepared and extracted before the next one in the sequence, (2) development can only take place towards one stope at a time, and (3) breaks in ore extraction are not permitted. Any constraints regarding stability were not considered. The order of operations for sequencing a simple panel and these constraints coincides with Figure 4-5. Recall that t_a is time to access a stope, t_e is time to completely extract a stope and t_p is time to prepare a stope.

1. Make a stope accessible in time $t_a \leq t_e$-t_p. The time constraint ensures constant flow of ore.
2. Prepare the stope in time t_p. Preparation includes:
    a. Adding required support.
    b. Drilling, loading, and blasting.
    c. Developing the extraction system.
3. Completely extract ore from the stope in time t_e.
4. Return to step 1.

Calculating access, extraction and preparation times require parameterization of operations as well as geometric information about the panel to be sequenced. Determining the net present value of a sequence also requires information about stope value, operating costs and discounting. Parameters used in the simple case include the following:

- Stope value and tonnage
- Stope vertex coordinates
- Annual discount rate in percent
- Stope preparation time in days
- Stope preparation cost in dollars per tonne
- Location of ore pass for haulage purposes
- Number of extraction units operating per stope
- Ore loading time and dumping time in seconds
- Haulage rate loaded and empty in meters per second
- Tonnes of ore hauled per load
- Tram equipment operating cost in dollars per day
- Development rate in meters per day
- Development cost in dollars per day

Value and tonnage of a stope would have been assessed prior to sequencing, possibly from implementing stope geometry optimization as described in Chapter 3. Only four vertices or corner-points describing each stope are used since panels are considered two dimensional in this case. Ore pass placement is described by a location along the strike direction shown in Figure 4-9. The type of extraction units are LHD's and several can be used to mine each stope. LHD's perform loading, hauling and dumping procedures which explains the need for only one equipment operating cost parameter. Development cost is the sum of all costs for labor, maintenance, material use and replacement (bits, drill steels, etc…), fuel or power, explosives, etc.

From these parameters, times and costs can be calculated for each operation required to mine a panel. Even though some operations such as stope extraction are continuous processes, a panel can be segmented up into a set of jobs. Based on constraints for the simple case, each job must be completed once it commences. This segmenting along with some of the parameters is shown in Figure 4-10. Optimization involves finding an order to execute these jobs in such a way that constraints are not violated and NPV is maximal.

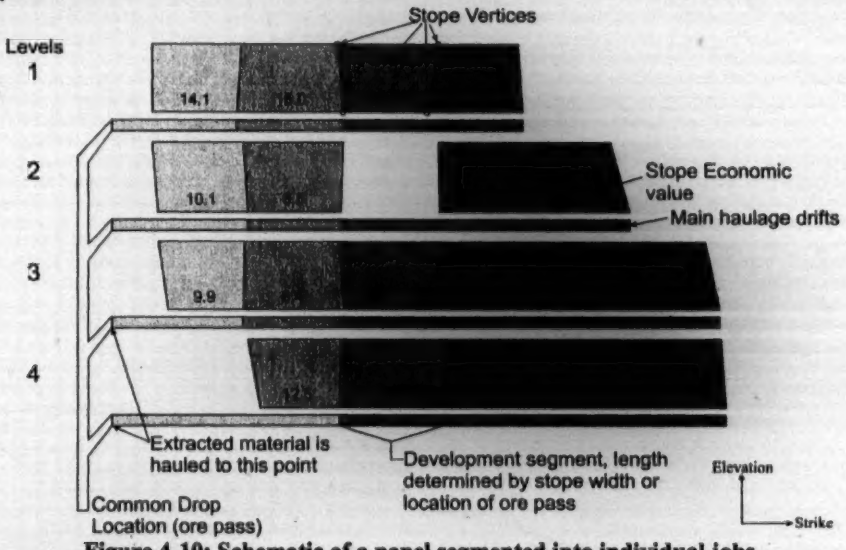

Figure 4-10: Schematic of a panel segmented into individual jobs.

4.4 Optimization

Two optimization algorithms were applied: simulated annealing and probabilistic decision making. These are only two approaches chosen to develop the framework for stope sequence optimization. Many other optimization algorithms could be applied including branch-and-bound, genetic algorithms, or exchange heuristics. Some of these methods are described in Chapter 2. SA and PDM were chosen to provide a random approach and a non-random approach to the problem.

Before discussing each optimization technique, the organization of information for a sequence will be covered. In the two dimensional case and assuming each stope is of similar dimensions, a panel can be represented by a matrix **S** with L rows and C columns. L is equivalent to the number of levels in the panel and C to the number of stopes along strike. Each entry in the stope matrix, $\mathbf{S}_{ij}$ i=1,…,L, j=1,…,C, stores stope properties and information including value, tonnage, neighboring stopes, mined status, etc. Another matrix **D**, which is the same size as **S**, stores development segments or drifts. Each entry, $\mathbf{D}_{ij}$ i=1,…,L, j=1,…,C, contains drift information such as mining cost, mining time, neighboring drifts, mining status, etc. Figure 4-11 shows these data structures. For the panel shown in Figure 4-10, **S** would be 4 by 6. Entries $\mathbf{S}_{15}$, $\mathbf{S}_{16}$, $\mathbf{S}_{23}$, $\mathbf{S}_{26}$, and $\mathbf{S}_{41}$ would be empty.

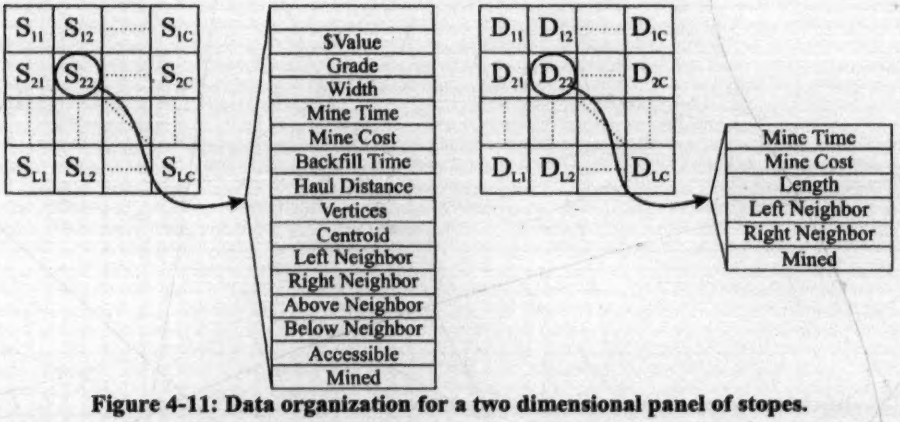

Figure 4-11: Data organization for a two dimensional panel of stopes.

In choosing a feasible stope sequence given the constraints considered for the simple sequencing case, a list of feasible stopes is built. For both optimization methods, the first decision encompasses all stopes; however, the remainder of decisions is based on the previous given the following;

For decision k

$\mathbf{S}_{ij}$ is feasible *iff* $t_{e,k-1} - t_{p,k} \geq t_{a,k}$

where $t_{e,k-1}$ is time to mine the stope at previous decision

$t_{p,k}$ is time to prepare the next stope

$t_{a,k}$ is time to make the next stope accessible

The time to make the next stope accessible can be calculated as follows:

$$t_{a,k} = \sum_{\mathbf{D}} t_{dm}(\mathbf{D}_{i(k)j}) \cdot \delta_{i(k)j}$$

where $t_{dm}(\cdot)$ is time to mine a drift $\cdot$

$\mathbf{D}_{i(k)j}$ is drift ij, and $i(k)$ means we only consider drifts on the same level as the stope for decision k

$\delta_{i(k)j}$ is an indicator: 1 if $\mathbf{D}_{i(k)j}$ is unmined, 0 otherwise

As a sequence is built by one of the optimization techniques, sequence order, completion time and profit matrices equivalent in size to **S** and **D** are filled. Completion time and profit matrices are used in calculating the net present value of the sequence. NPV is calculated as with the following equation.

$$NPV = \sum_{i=1}^{m}\sum_{j=1}^{n}\left[\mathbf{SP}_{ij}\cdot\left(1+\tfrac{r}{365}\right)^{-\mathbf{ST}_{ij}} + \mathbf{DP}_{ij}\cdot\left(1+\tfrac{r}{365}\right)^{-\mathbf{DT}_{ij}}\right]$$

where **SP** is the stope profit matrix

ST is the time a stope is completed in days

DP is the drift profit matrix

DT is the time a drift is completed

r is the annual discount rate

It should be noted that anywhere in the panel a stope does not exist, the profit and time matrices just store a value of zero so the NPV calculation is not effected.

4.4.1 Optimization with Simulated Annealing

Simulated annealing was chosen mainly for its ease of implementation. Having a simple algorithm available was ideal for moving forward with developing a framework for stope sequence optimization. Unlike stope geometry optimization as discussed in Chapter 3, stope sequence optimization is a combinatorial problem. We are randomly perturbing integers to hone in on a stope sequence offering maximum NPV. The integer scale ranges from 1 to the total number of stopes to be sequenced.

Simulated annealing is easily adapted to combinatorial problems. Random numbers are drawn from a uniform distribution and range from zero to one. Drawn values are scaled to range from one to the number of stopes and then truncated to integers. This allows the algorithm to randomly draw a sequential order for the stopes to be mined in. There are more details to this random selection that prevent infeasible sequences from being chosen.

The basic process for simulated annealing as discussed in Chapter 2 is followed: (1) defined initial sequence parameters and constraints, (2) randomly choose a feasible sequence and calculate its NPV, (3) randomly select two stopes and swap their order in the sequence, (4) evaluate the NPV of the new sequence and accept or reject it, and (5) reduce the annealing temperature accordingly and return to step 3. This process is shown graphically by Figure 4-12. The sequence parameterization considered for Step 1 is defined in Section 4.3.

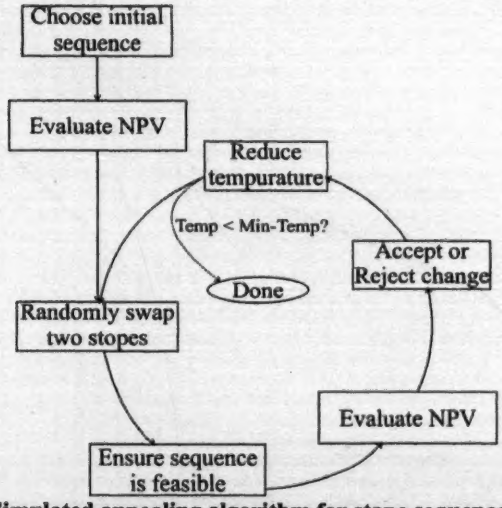

Figure 4-12: Simulated annealing algorithm for stope sequence optimization.

Choosing an initial sequence to start the optimization algorithm off with is random. For the scenario developed here and for a given panel, any of the stopes can be the first stope mined. Once the panel is started, we only need to select stopes from those that are feasible given the previous one selected. A list of feasible stopes is generated using the time constraints as defined in Section 4.4. Stopes in the list are assigned temporary indices ranging from one to the number in the list. One of these indices is then randomly drawn to select the next stope in the sequence. In the event that a sequence eventually reaches a dead end (not all stopes can be feasibly mined), the sequence is deleted and the process is started over. The optimization process is not started until an initial feasible sequence is found. Initial annealing temperature is set as the NPV of this sequence.

With an initial feasible sequence selected, optimization can begin. Small changes to the sequence are made by exchanging the order of two stopes per iteration. There were two ways in which a pair of stopes to swap could be selected. One involves choosing the first randomly, and then generating a list of those it could feasibly be swapped with from which the second is selected. The second method involves choosing both stopes randomly without being concerned with feasibility, then mining the sequence as close as possible to that proposed by the swap. This method was used in the algorithm ('Ensure sequence is feasible' stage in Figure 4-12). Determining a list of stopes that can be feasibly swapped is time consuming. In most cases, random selection of both stopes not necessarily being feasible leads to a new sequence being evaluated. The second method allows more sequences to be evaluated in less time.

Acceptance or rejection of new sequences is done based on probability. If the new sequence has a higher NPV or may lead to a better sequence, it is accepted. Since only two stopes are swapped at a time, it is important to occasionally accept sequences with lesser NPV. It may take several sequential swaps to yield a sequence with higher NPV than the current maximum. A simple example of this can be shown with a hypothetical panel of nine stopes, see Figure 4-13.

The simulated annealing temperature is reduced differently for accepted or rejected sequence proposals. If a sequence is accepted, the temperature is reduced more aggressively than if it is rejected. Optimization is terminated when the temperature decreases below a user specified minimum. Since the temperature represents a dollar value, it can be selected based on the level of improvement in a sequence a user is concerned with.

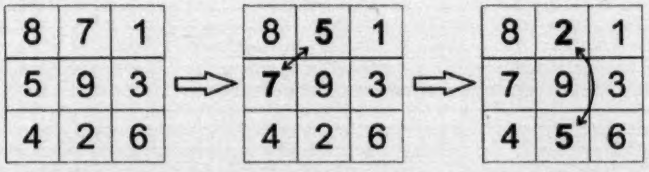

Figure 4-13: Process of accepting suboptimal sequences to reach more optimal solutions.

Connectedness of Feasible Sets

Because the simulated annealing algorithm described above only exchanges two stopes per iteration, the connectedness of feasible sets must be discussed. It is likely that two feasible sets cannot be linked by a set of exchange operations, making them disconnected. If from some random feasible sequence we cannot visit all possible sequences, the set is disconnected. In the case, a random restart component can be added to the algorithm.

Random restart could be implemented as an additional loop to the simulated annealing optimization algorithm. When optimization is finished in Figure 4-12, another initial sequence would be chosen and optimization would restart. How many restarts are carried out is a subjective matter. This value can be set discretely or it may depend on improvement over time calculations or another temperature parameter.

4.4.2 Optimization with Probabilistic Decision Making

Use of a logic driven decision making process was implemented because there are logical reasons why mining one stope prior to another provides a higher NPV. Simulated annealing chooses a sequence randomly whereas the algorithm to be developed in this section will choose a sequence based on stope and operations characteristics. The algorithm is a decision making heuristic that attempts to choose the next stope in a sequence that may lead to an optimal solution.

PDM was developed for the sequencing scenario discussed in Section 4.3. Making decisions is based on the properties of each stope in the sequence. Those that were selected aid in sequencing because they make a stope seem more or less desirable in the selection process. They include:

1. Revenue from the stope (value less preparation costs)
2. Time to mine the stope
3. Time to make the stope accessible (based on which drifts must be mined to reach it)
4. Cost to access the stope (Expenses incurred to make the stope accessible)

For the scenario this algorithm was developed for, there are only four properties. In a more complex mining environment, there may be many more that come into the decision making process. These properties are used in calculating a probability for each stope, which reflects how probable the stope is to lead to an optimal solution.

Referring to the stope matrix in Figure 4-11, each stope is assigned a vector of weights and properties. The vector of properties has four elements containing the value of each property explained above. Revenue is calculated as the stopes value less any preparation costs. Time to mine the stope depends on its tonnage, the haulage routes, and equipment used to mine it. Time to access the stope is based on access drifts and equipment used to mine them. Costs to access the stope are based on drifts to be mined and operating costs to mine them. Weights are assigned to each of these properties based on there impact on NPV of the panel. Revenue is positive so its initial weight is positive. The remaining three are negative since more time means more discounting and higher costs take away from the revenue. The initial weight vector was {-0.5, 0.5, 0.5, 0.5}.

Equation 4-1 is used for calculating probabilities. Since exponentials are used, stope properties had to be normalized by the maximum observed value for each. Otherwise, the equation becomes unstable.

$$\mathbf{P}_{ij} = \frac{\exp\left(\mathbf{w}_{ij}^T \mathbf{F}_{ij}\right)}{\sum_{\mathbf{S}} \exp\left(\mathbf{w}_{ij}^T \mathbf{F}_{ij}\right)} \qquad 4\text{-}1$$

Where $\mathbf{P}_{ij}$ is the probability of stope *ij* being the next best decision and $\mathbf{w}_{ij}$ and $\mathbf{F}_{ij}$ are the weight and property vectors for stope *ij*. **w** and **F** each have four elements. $\mathbf{w}_{ijk}$ is the weight applied to $\mathbf{F}_{ijk}$, k=1,...4 (k=1 is stope revenue, k=2 is mining time, k=3 is accessing time and k=4 is accessing cost).

For each state of the panel, the stope with the highest probability is chosen to be next in the sequence. Of course, this algorithm would only provide one sequence based on the initial weights if there was no way of updating the weights. Properties may require a different weighting to lead to an optimal solution. Weights are updated using a steepest descent algorithm. When a decision leads to a lower NPV it is considered poor. The weights for that stope are updated along the negative gradient of Equation 5-1 such that it is not chosen next in the sequence.

Steepest descent works by computing the steepest descent direction (the gradient), then using a line search algorithm to step to the minimum along that direction [2]. Rather than stepping to a minimum however, we would like to step far enough so that the highest probability is slightly less than the second highest. Weights are not changed more than

required to cause a change in the sequence. The gradient of Equation 4-1 is shown by Equation 4-2. The result is a vector the same size as the weight and property vectors. Weights are updated along its negative. A backtracking line-search was used [2].

$$\frac{\partial \mathbf{P}_{ij}}{\partial \mathbf{w}_{ijk}} = \frac{\left(\sum \mathbf{P} - \mathbf{P}_{ij}\right) \cdot \mathbf{F}_{ijk} \cdot \exp\left(\mathbf{w}_{ijk} \cdot \mathbf{F}_{ijk}\right)}{\left[\sum \mathbf{P}\right]^2} \qquad 4\text{-}2$$

Where k is the weight and property of interest (k=1,…,4), $\mathbf{w}_{ijk}$ and $\mathbf{F}_{ijk}$ are weight k and property k for stope ij and $\mathbf{P}$ is the matrix of probabilities for all stopes.

Determining if a decision is poor (leads to a lower NPV) is done by attempting to calculate an upper bound on NPV. Bounds are calculated by splicing a current sequence decision into the previous best complete sequence. This requires an initial sequence however. A greedy algorithm was used to do this [6]. Greedy algorithms work by making the best choice at the current state of a problem. The choice is locally optimal. For stope sequencing, a greedy algorithm was developed using discounted revenue as a parameter. This is based on attaining profit as soon as possible so that it undergoes a lower degree of discounting. The algorithm is summarized below.

1. Create a list of feasible stopes.
2. Choose the stope with highest discounted revenue.
3. Repeat 1 and 2 until the full sequence is complete.

The greedy algorithm, equation for probabilities and steepest descent are combined to form the PDM algorithm, which is described by Figure 4-14.

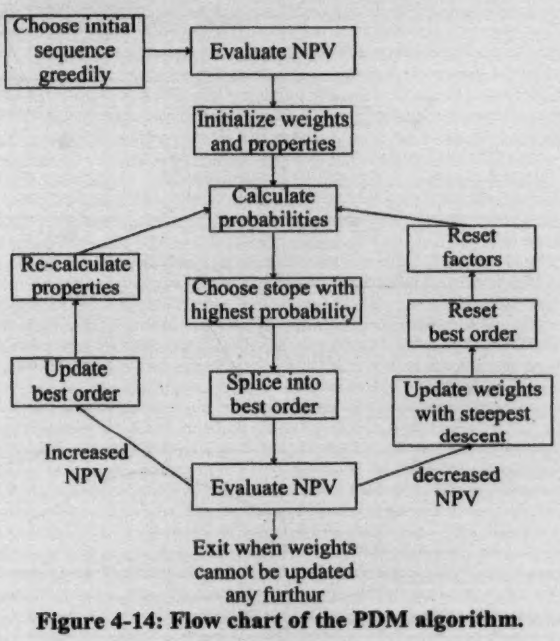

Figure 4-14: Flow chart of the PDM algorithm.

A step that requires more detail is splicing a decision into the previous best sequence. As the algorithm is working on a sequence, what is available is the start of a new sequence and the current best sequence. These are combined to generate a full sequence. The new sequence is followed up to and including the current decision, then remaining decisions follow the previous best order. In this way, an approximate NPV for a particular decision can be calculated.

Rather than basing stopping on an annealing temperature like stope geometry optimization, we now have a gradient. Eventually, the weight vectors for each stope will settle into a local minimum for their position in the sequence. At this point the gradient will be extremely small or zero and no updating can be done. Since the sequence cannot be changed optimization is complete.

Both algorithms described for stope sequencing were tested and compared with a panel of stopes. There were 70 stopes in the panel and they were mined under the sequencing framework developed. A detailed description of the panel and sequencing results are provided in Chapter 5.

Chapter 5 Case Studies

This chapter focuses on applying the optimization algorithms developed in Chapters 3 and 4 to different mining scenarios. As stope geometry optimization is further developed than stope sequencing, more cases are discussed. Three scenarios were developed for stope geometry: (1) a very simple two dimensional case such that the objective surface can be observed, (2) a large stope designed for a sublevel stope mine, and (3) a small stope for cut-and-fill mining. These mining methods are explained in Chapter 1. Recall that the objective surface at a particular state is the value of a stope. The state is that stope's current design or vertex configuration. Objectives of these case studies are to show that a stope's value can be improved with the algorithm presented in Chapter 3 and to identify areas for improvement.

Only one study was carried out for stope sequence optimization. There were several objectives: show that sequence value can be improved with both algorithms presented in Chapter 4, compare both algorithms, and identify any issues with the framework that require improvement. A more detailed discussion of future work for both optimization strategies will be given in the next Chapter.

5.1 Application: Stope Geometry Optimization

Three scenarios were put together to test the stope geometry optimization algorithm. The first was created not to test optimization but to visualize some of the many possible objective surfaces that are being searched by the optimization algorithm. By doing this, we can see how complex the surface can be and why other optimization techniques may not find an optimal solution. The second and third scenarios were developed to test various aspects of the algorithm including run time and the effect of block model clipping accuracy on solutions. These two cases also allow a comparison when dealing with very different stope designs. A more detailed description of these scenarios follows. A graphical representation is provided in Figure 5-1.

1. This is a very simple two dimensional case. It is assumed that the block model consists only of one row of blocks. The stope consists of two ends that are positioned along the row of blocks. Its value is the sum of the blocks between the ends.
2. Sublevel stope mining was targeted for this case. The stope is large relative to the block sizes in the economic model. It consists of 18 vertices and various constraints were imposed on them.
3. This stope was designed for a cut-and-fill operation. It is essentially a development drift placed in ore bearing rock. 32 vertices make up its design and it is small relative to the blocks in the economic model.

For the second and third cases, a hypothetical economic model was generated. Blocks were sized to be 1 meter cubes and the model was 40 blocks in each direction. Economic values range from -200 to 946. The average value for the entire model is -107. Figure 5-2 shows the distribution not including the -200 values as most of these blocks are away from the ore bearing zone. An image of the model is shown in Figure 5-3.

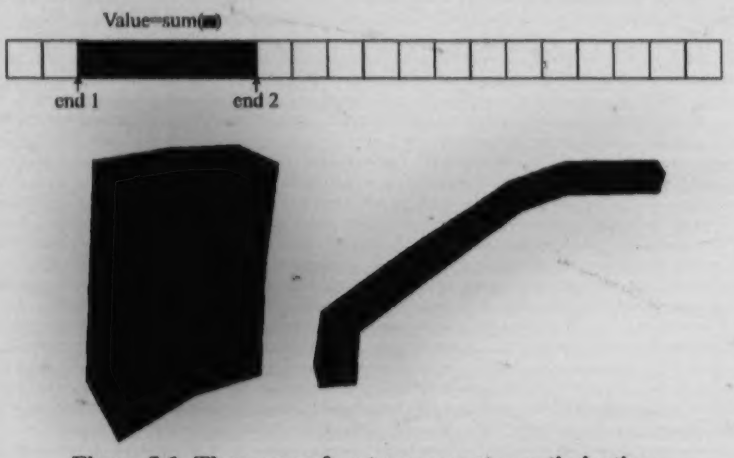

Figure 5-1: Three cases for stope geometry optimization. Case 1 (top), case 2 (bottom left), case 3 (bottom right).

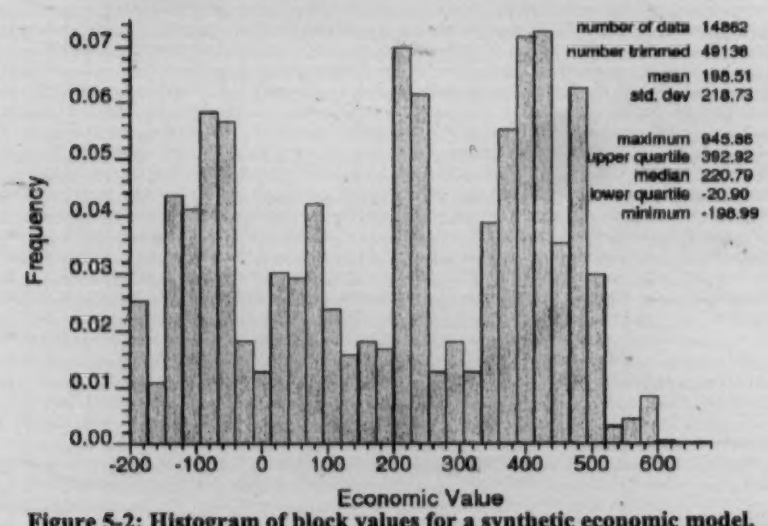

Figure 5-2: Histogram of block values for a synthetic economic model.

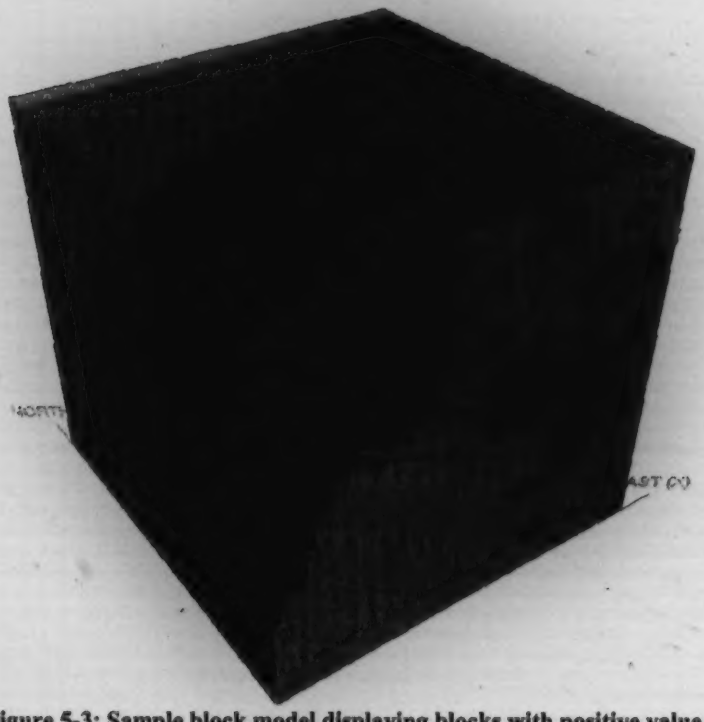

Figure 5-3: Sample block model displaying blocks with positive value.

5.1.1 Case 1: The Objective Surface

Since we can only visualize in three dimensions, this case is not feasible for actual mining applications. It is however a case where the objective surface can be plotted. Consider an economic block model in three dimensions that is one block in the y and z directions and several blocks in the x direction. Assume that the stope is a simple tunnel that runs in the x direction. It has two ends and its cross section is consistent with the economic blocks in the yz plane. The value of this stope is the sum of all blocks between its ends, see Figure 5-4.

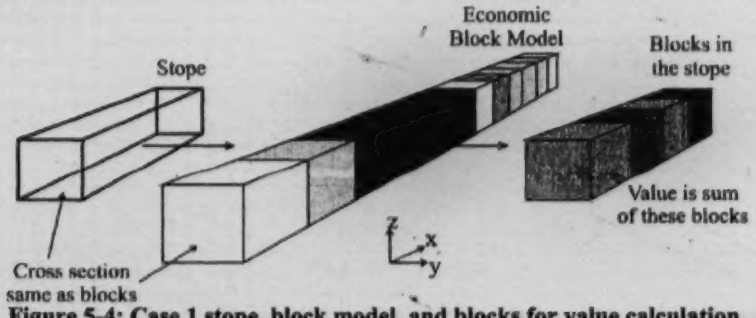

Figure 5-4: Case 1 stope, block model, and blocks for value calculation.

Optimization of the stope for this case would involve shifting the ends of the stope in the x direction and evaluating its value until the maximum is found. The objective surface for this problem is three dimensional. The positions of the ends will form two orthogonal axes and the stope value will form the third axis. Generating an image of this surface was accomplished by selecting a set of locations for each endpoint and iterating though them calculating the value for each. The surface is symmetric about the line formed where the positions of the two ends are identical. Three synthetic block models were constructed and surfaces where generated for each, see Figure 5-5. Distributions of values along the block models were as follows

Surface 1: A mathematical function was used: $f(x) = 3x^5 - 25x^3 + 60x$.

Surface 2: A row of blocks was extracted from the economic model generated for cases 2 and 3.

Surface 3: A random vector was created using uniformly distributed numbers between -50 and 50.

Even with this simple case, the state space to search for a maximum value can be complex. An optimization algorithm that is capable of escaping local maxima is needed if a global maximum is to be found. Gradient based optimization techniques such as steepest ascent and Newton's method are likely to converge on a local maximum nearest the initial stope design or state. Going beyond this simple case to a stope with eight vertices for example and a large three dimensional economic block model increases the dimensionality of the problem significantly. If the stope vertices are shifted along only one axis, the objective surface would exist in eight dimensions. If vertices are free to be shifted in all directions, we are dealing with 24 dimensions. Given the complexity of the economic model, the objective surface will be highly convoluted.

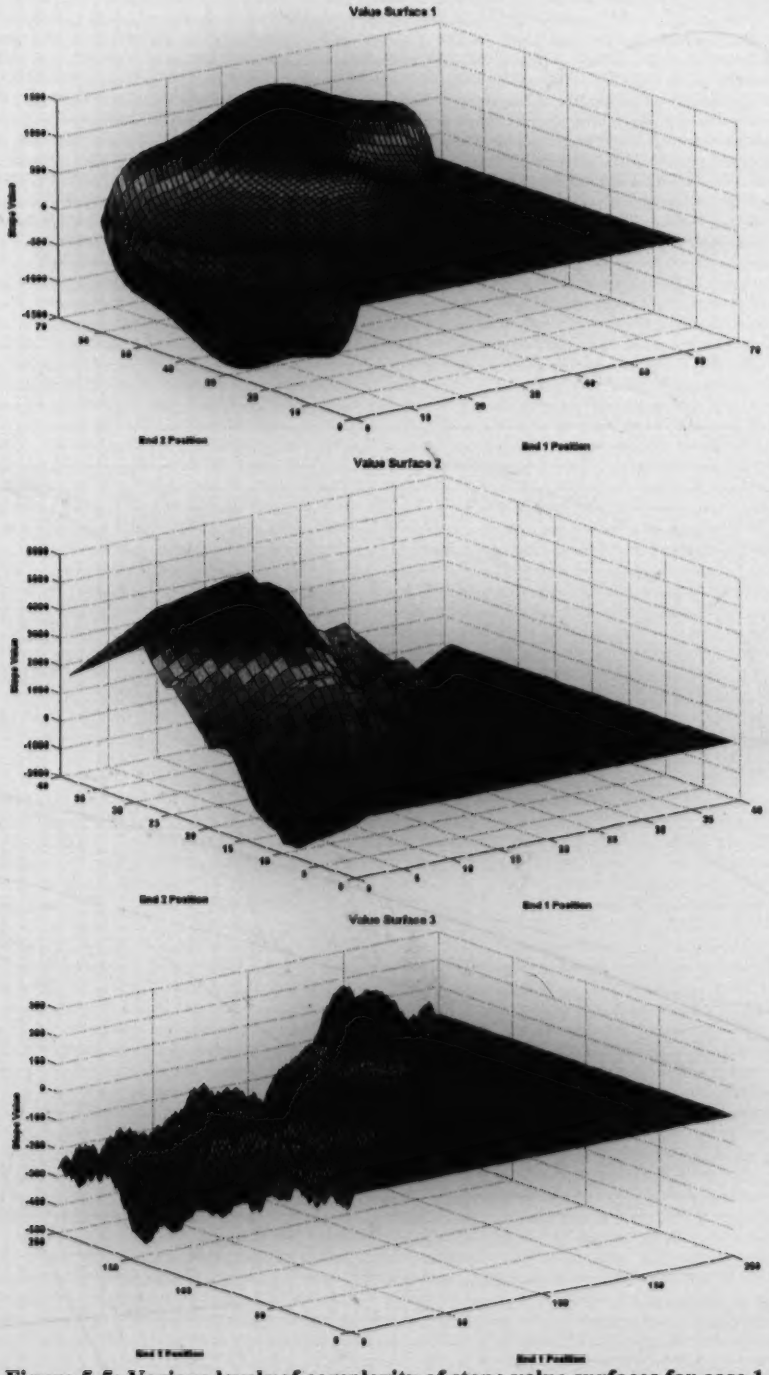

Figure 5-5: Various levels of complexity of stope value surfaces for case 1.

5.1.2 Case 2: Sublevel Stoping

A stope for a sublevel stoping scenario was designed using the synthetic economic model. Before designing the stope, several assumptions about the underlying mine design must be made: (1) the mine design has fixed levels along which development and haulage drifts are extracted, (2) fan drilling takes place from the fixed levels, (3) drilling is done from above and below each stope allowing for some flexibility in terms of stope design. With these in consideration, the stope for this case was designed using three plan views of the economic model, see Figure 5-6. Two of the views coincide with the upper and lower levels and the third with an intermediate level between them. Six vertices had to be positioned on each view.

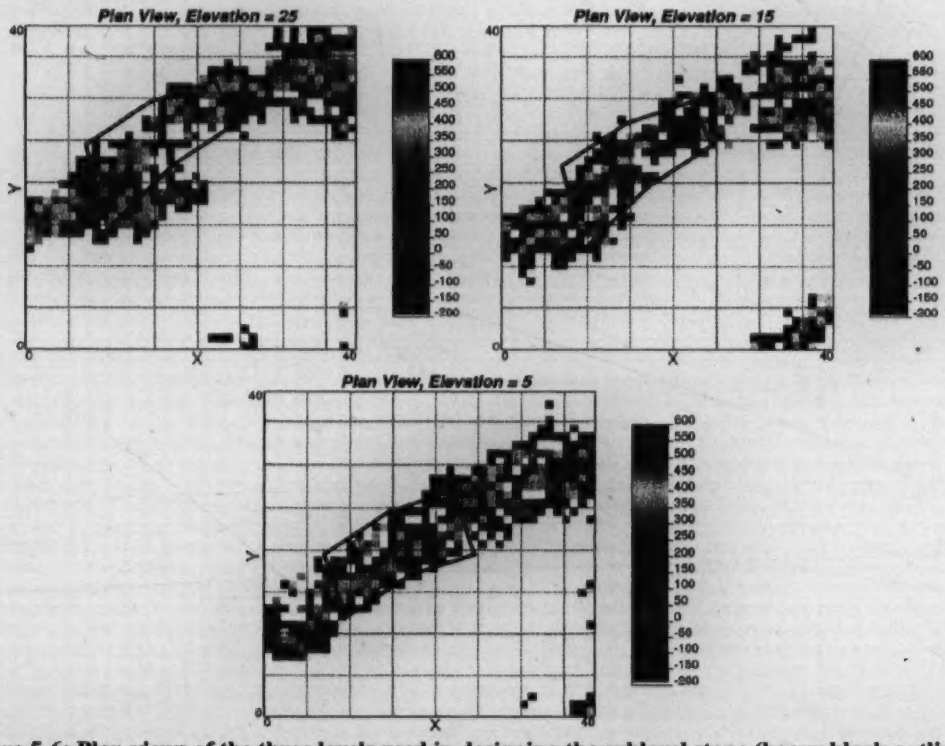

Figure 5-6: Plan views of the three levels used in designing the sublevel stope (heavy black outlines).

Vertex movement constraints were also considered for this problem. It was assumed that the influence of neighboring stopes prevent any flexibility along the sides of the stope. Side vertices were locked in their initial positions. Because the mine design has fixed levels, central-upper and central-lower vertices were restricted to movement in the aerial plane. The two remaining central vertices were permitted to move in any direction. Use of fan drilling from above and below should give enough flexibility to do this. These constraints are shown graphically in Figure 5-7. To provide a more achievable optimized stope in terms of drilling and blasting, corner and face deviation constraints were imposed. The acute angle of corners could not be less than 60 degrees and the obtuse angle across faces could not exceed 210 degrees.

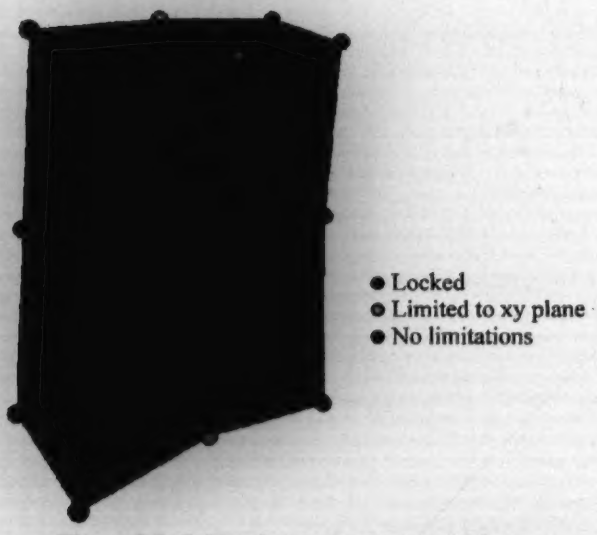

Figure 5-7: Sublevel stope vertex constraints.

Optimization was carried out twice: once with no block refinement and again with blocks refined. This was done to compare run-time and solution differences. Table 5-1 summarizes the results. Commenting on run-time first, it seems that the data structures for dealing with partial blocks require improvement. Differences in solutions between using refinement and not refining blocks are small. This can be explained by the size of the stope relative to the blocks in the economic model. The stope is very large in this regard. A small percentage of the total blocks inside the stope are found on the stope's perimeter.

Table 5-1: Sublevel stope optimization summary.

Block Refinement	Run-time	Initial Value	Final Value	% Improvement
NO	0 min 12 sec	271,898	314,002	15.5
YES	8 min 48 sec	267,718	306,340	14.4

An improvement of 15 % is quite substantial. This is especially true since the stope design was significantly constrained. If we relax the problem and let all of the vertices that were previously locked move in the aerial plane, values exceeding 580,000 can be achieved. Taking a look at the initial and final designs in Figure 5-8, we can see that the stope has not undergone much modification. It is still realistic in terms of drilling and blasting capabilities. Regarding areas that require improvement, one has been identified by this case and that is the inefficiency of the block refinement procedure.

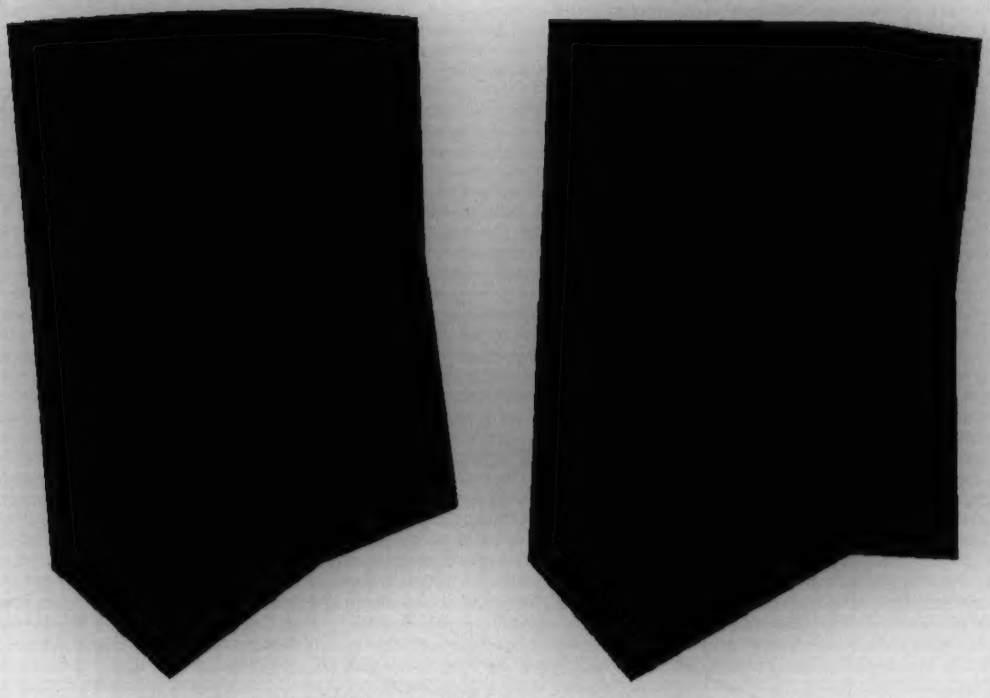

Figure 5-8: Initial (left) and optimized (right) sublevel stope designs.

5.1.3 Case 3: Cut-and-Fill

Cut-and-fill mining is a completely different environment than sublevel stoping. In this case, a stope is similar to a development or haulage drift. It is advanced into ore for production purposes. Some important aspects about the underlying mine design that impact stope design are: (1) there are a fixed set of levels within which stopes or drifts are mined, (2) the levels are separated by a vertical distance that is equal to the height of the stopes, and (3) drifts are designed to have a specific cross-sectional shape for stability purposes. For case 3, the stope was designed with a square cross-section having a width and height of 2.5 meters. The stope was designed with one plan view, see Figure 5-9. A total of 32 vertices and 60 faces describe the stope.

Constraints for this problem consisted of vertex, mining width and deviation constraints. A segment of the stope was present for access. This was assumed a fixed component of the mine design so those vertices were locked. The remaining vertices were permitted to move in the aerial plane only. This keeps the drift height fixed. Mining width was limited to range between 2.25 and 2.75 meters. Corner deviations were kept very tight such that the square cross section is maintained fairly closely. Face deviations, which essentially amount to bends along the length of the drift, were permitted to range from 180 to 270 degrees for obtuse angles.

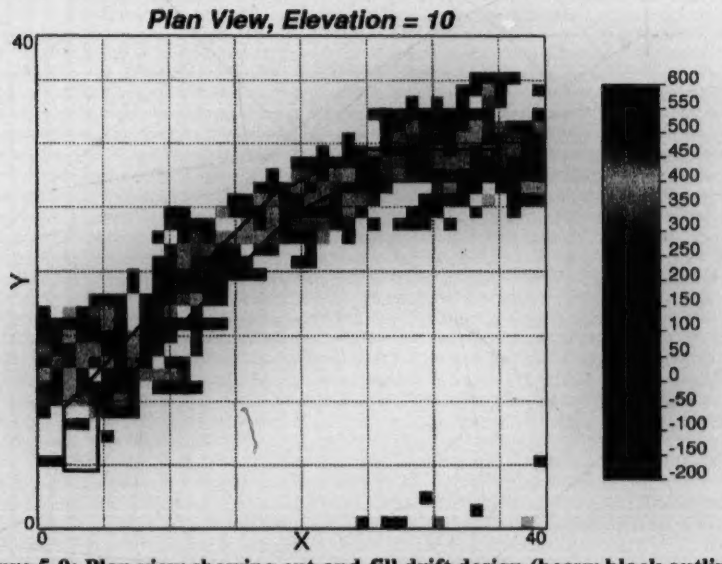

Figure 5-9: Plan view showing cut-and-fill drift design (heavy black outline).

Two optimization runs were carried out: one with block refinement and one without. The main focus of these runs was to show that the algorithm is flexible in terms of stope design and constraints. Results are summarized in Table 5-2. Block refinement was important in this case. The stope was small relative to the size of the blocks in the economic model. When no refinement was considered, the set of blocks within the stope form a poor representation of the stope. Figure 5-10 compares the set of blocks inside the stope for both optimization runs. Initial and final stope designs are shown in Figure 5-11. They are quite similar.

Table 5-2: Cut-and-fill stope optimization summary.

Block Refinement	Run-time	Initial Value	Final Value	% Improvement
NO	N/A	38,104	52,288	37.2
YES	N/A	31,648	41,963	32.6

Improvement of the initial design was substantial at 32.6 %. However, an area that requires improvement in the algorithm is likely responsible for generating too large an increase. A mining width range of 2.25 to 2.75 meters was mentioned above, but to preserve the cross-section of a stope the width should be 2.5 meters. If the width were constrained to 2.5 meters, then all vertex moves would violate the constraint and be rejected. To combat this problem, the user should be able to identify subsets of vertices that can be moved in tandem with one movement vector. Moving a set of vertices that form a cross-section of the stope together would maintain a width constraint.

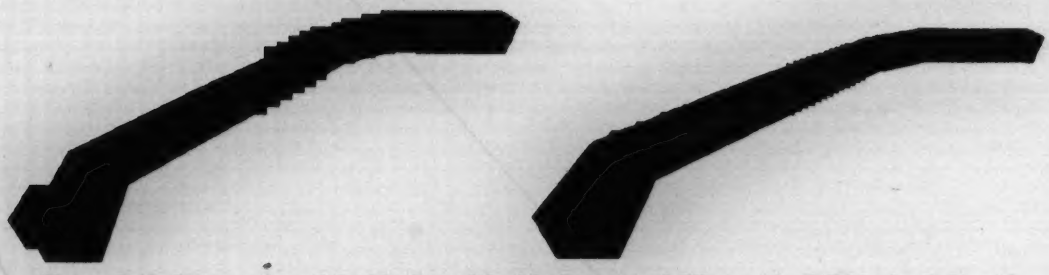

Figure 5-10: Intersection of the block model with the cut-and-fill stope by full blocks (left) and refined blocks (right).

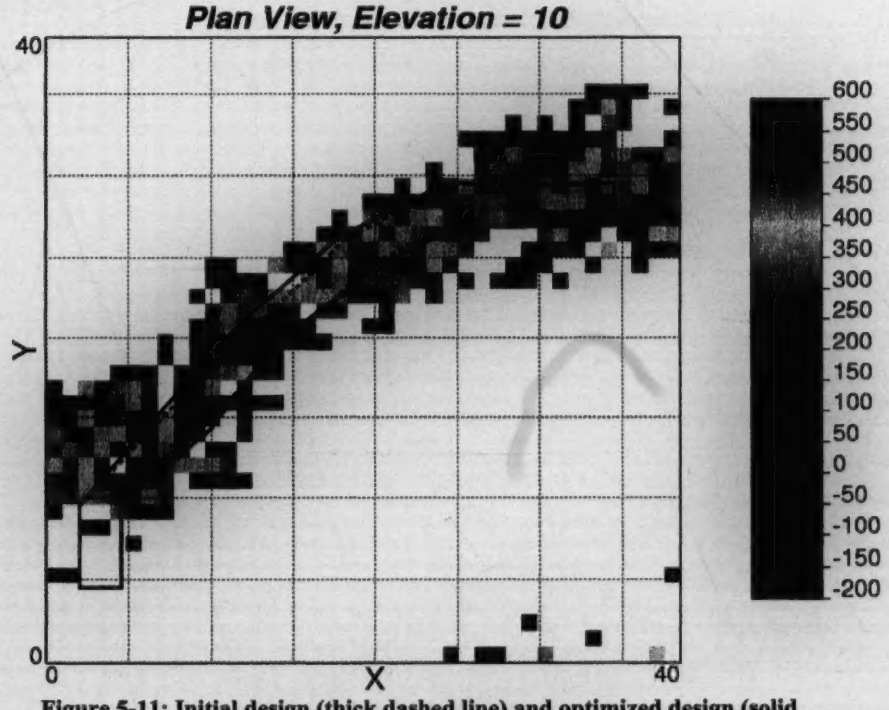

Figure 5-11: Initial design (thick dashed line) and optimized design (solid black line) for the cut-and-fill stope.

5.2 Application: Stope Sequence Optimization

Only one set of stopes will be processed by the stope sequencing algorithms presented in Chapter 4. It will be shown that the optimization of sequences under the framework developed does result in an increase in net present value. Improvement is relative to the value of the first sequence chosen. Recall that the first sequence is chosen randomly by

simulated annealing and by use of a greedy algorithm for probabilistic decision making. Optimization algorithms will be compared using several results: value of the initial sequence, value of the final sequence, run-time, and sequence order differences.

For this application, the set of stopes were designed and used in an actual mining operation. The panel comprises 70 stopes in a sublevel stoping environment, see Figure 5-12. Since the framework developed in Chapter 4 is limited to one set of mining procedures, the actual methods used to mine the panel could not be simulated. Mining procedures for the framework were used. Recall that one development crew and one mining crew are used. Constraints are: one stope must be undergoing extraction at all times, jobs that are started must be fully complete before that crew moves on, and development can only take place towards one stope at a time.

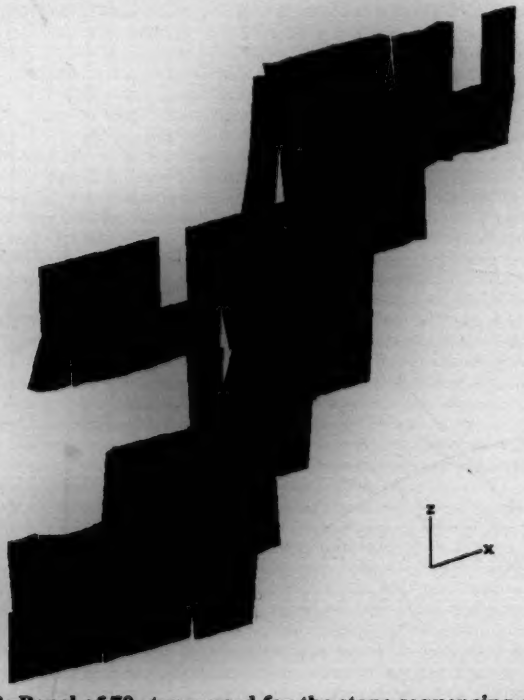

Figure 5-12: Panel of 70 stopes used for the stope sequencing application.

The panel shown in Figure 5-12 is spread over eight levels. Its geometry allows us to project the panel into two dimensions so that the sequencing framework can be applied. Figure 5-13 shows the 2D projection with stopes color-coded by their economic value. Values range from $ 90,000 to more than $ 800,000. Because the mining scenario is simple with few constraints, we would expect sequencing to mine more valuable stopes earlier on in the sequence. They will be discounted less resulting in a higher NPV.

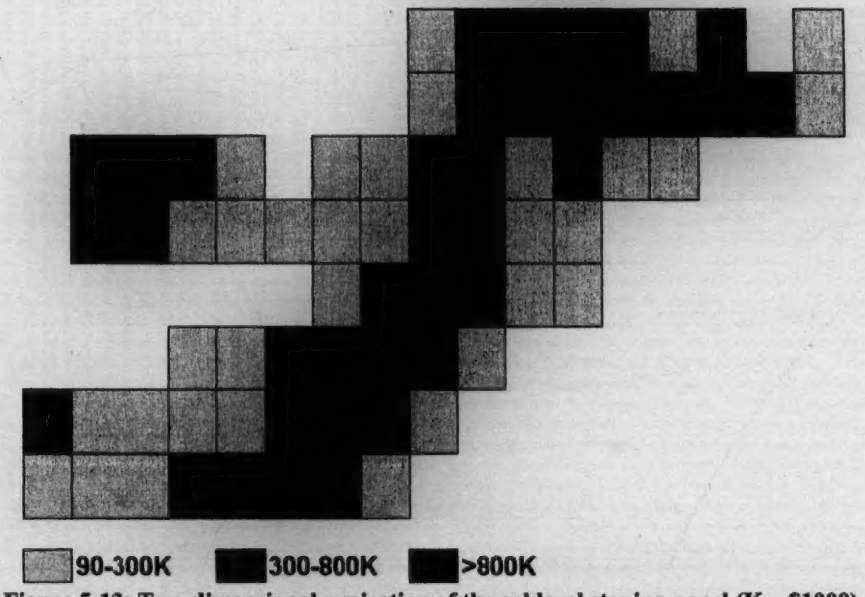

Figure 5-13: Two-dimensional projection of the sublevel stoping panel (K = $1000).

Optimization results are summarized in Table 5-3 for initial value, final value and run-time. Looking at initial sequence values, that for simulated annealing is sub-optimal when compared to that for PDM. The first sequence for simulated annealing is chosen randomly and that for PDM is chosen using a greedy algorithm. The greedy algorithm makes sequence decisions based only on stope revenue. This has advantages over randomly choosing an initial sequence: it will provide consistent results and it is based on some logic in terms of calculating NPV. The first sequence chosen by simulated annealing will not produce consistent results as it is random.

Table 5-3: Stope sequence optimization results for simulated annealing and PDM. Values are in millions of dollars and times in seconds.

Stopes	Simulated Annealing			PDM			ΔNPV_f	ΔT
	NPV_i	NPV_f	Time	NPV_i	NPV_f	Time		
70	11.694	12.064	16.03	11.952	12.004	88.61	0.061	-72.58

Final NPV values show better results using simulated annealing. This could be for several reasons. One is the simplicity of the problem. There are very few constraints. Another reason has to do with the PDM algorithm. An initial sequence is chosen and then used in calculating NPV for individual sequence decisions. PDM makes sequence decision one stope at a time. If the initial sequence requires multiple stopes be rearranged to improve NPV the algorithm will never make an improvement.

Comparing run-time, simulated annealing ran substantially faster. Calculation of gradients and performing a line search are slow operations. During the time to carry out those operations, simulated annealing could have evaluated several new sequences. For PDM, a large amount of time, from 23 seconds to completion, was spent with no improvements being made. At this point, no individual decisions were being accepted and the weights were being updated to some local minimum for each stope. Runs are characterized in Figure 5-14.

A comparison of sequence orders can also be made. Figure 5-15 shows the sequence for each algorithm as well as the differences. Both algorithms tended to mine more valuable stopes earlier in the sequence. A notable difference can be made for the valuable stopes in the top two levels of the panel. Simulated annealing mined these stopes prior to 21st in the sequence whereas PDM mined them later than 23rd and prior to 58th. These stopes would be discounted more for PDM, which likely contributes to some of the difference in final NPV between the two algorithms.

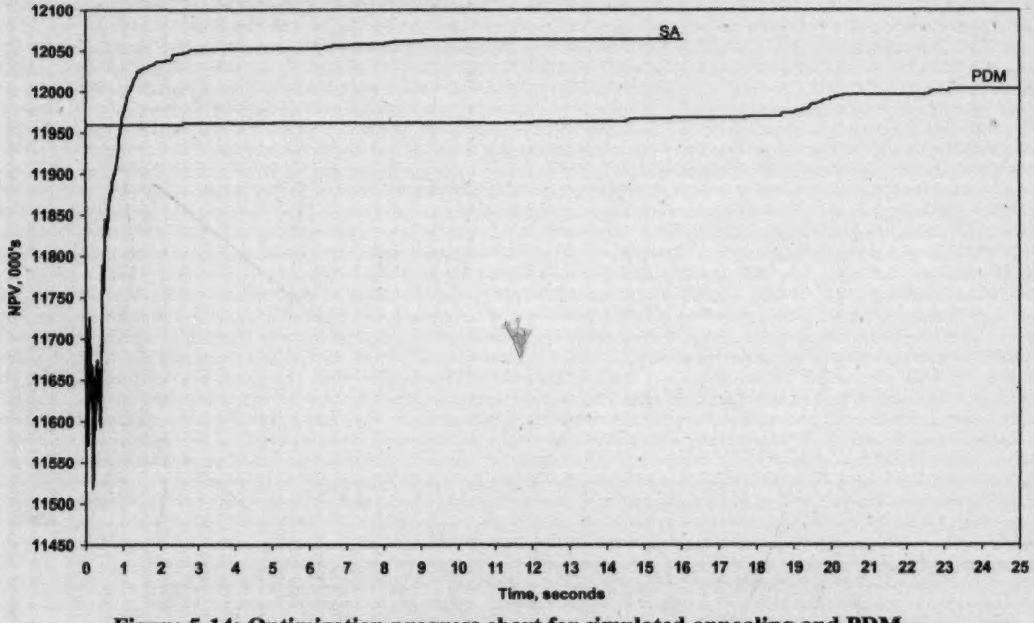

Figure 5-14: Optimization progress chart for simulated annealing and PDM.

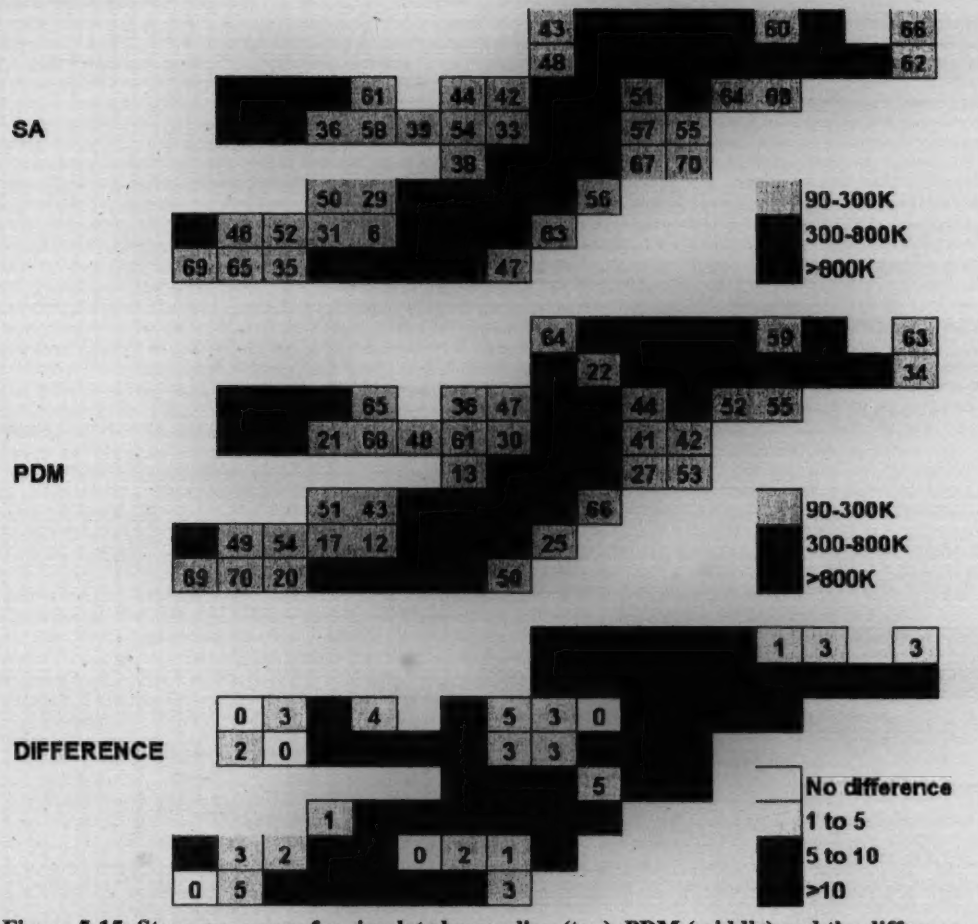

Figure 5-15: Stope sequences for simulated annealing (top), PDM (middle) and the difference between them (bottom).

An additional set of results was generated for simulated annealing to explore the need for a random restart. Fifteen runs were carried out, each with a different random number seed. Several features were looked at including the average sequence, standard deviation of each stope's position in the sequence and some basic statistics on the NPV results. Basic statistics are given in Table 5-4. Sequence aspects are shown graphically in figure 5-16. Based on NPV responses along, use of a random restart component would be beneficial. From minimum to maximum, a gain of 21,904 dollars could be realized. The standard deviation in NPV is quite low considering we are dealing with millions of dollars. Mean and standard deviation of the resulting sequences were looked at, see Figure 5-16. On average, more valuable stopes were chosen earlier in the sequence. Variability in sequence order for individual stopes is lower for stopes chosen near the start and end of the sequence and higher for those in between. Due to the nature of NPV calculations, it was expected that order variance would increase towards the end of the sequence. Incurring the most discounting, stopes at the end of the sequence have little effect on NPV.

Table 5-4: SA NPV Statistics

Statistic	Value ($)
Minimum	12,043,998.00
Maximum	12,065,902.00
Mean	12,055,929.70
Standard Deviation	5,705.08

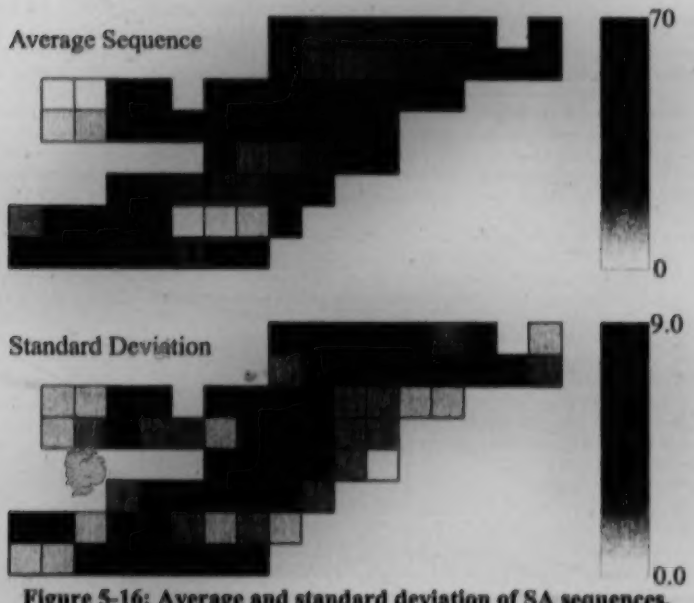

Figure 5-16: Average and standard deviation of SA sequences.

This study has helped identify areas for improvement for the stope sequencing algorithms, specifically PDM. Simulated annealing was straightforward to apply and will not be changed. It will have to be adapted to other mining scenarios that differ from that used in the framework. The addition of a random restart component would be beneficial. Regarding PDM, a noted problem is that it learns off an initial sequence and from a specified set of initial weights. Learning is halted when the weights settle into a local minimum for each stope. Other optimization algorithms for updating the weights that can escape these minima may improve results. Perturbation methods such as random restart, penalization and tunneling could be implemented [12]. When a local minimum is reached, random restart would simply choose a new weight vector for that stope. Penalization methods would add some "bump" to the function where the local minimum occurred. The steepest descent algorithm could then descend along the new gradient, perhaps towards a better local minimum. Tunneling searches from a local minimum for a new starting point. Any of these may improve the PDM algorithm. This was one of many methods that could be applied to the problem of stope sequence optimization. Other techniques should be implemented as well.

Chapter 6 Conclusions and Future Work

Economics determine if a mining project will be profitable. Feasibility studies are carried out to describe virtually all aspects of a project such that a bottom line value can be determined. Many components of the study are accompanied by uncertainty, which works through to the bottom line as well. Two components that have uncertainty are numerical modeling of an ore deposit and the mine design proposed to exploit the resource. The mine design will include parameters such as operating and processing costs. These along with the commodity price and deposit property models are used to generate an economic model of the deposit. The viability of a particular mine design can be evaluated using the economic model. For any project, the mine design should be brought to its full economic potential given the uncertainty involved to determine if it is feasible.

This thesis has presented two ways that the economic potential of an underground mine design can be improved. These components were stope designs and stope sequences. It would be very unlikely that initial stope designs or sequences for a project offer their maximum potential. However, they can be optimized increase value and in such a way that uncertainty is accounted for. Uncertainty exists in the numerical models (rock type and ore grade), in economic parameters (commodity price and operating and processing costs) and in processing recovery. These are all used in generating a set of possibilities for an economic model. These possibilities can be combined into an expected economic model. If optimization of stope designs and sequences are done based on this model, they provide their maximum potential in expected value. Thus uncertainty is accounted for.

The main contribution made by this thesis is that numerical optimization techniques can be used in an underground mining context to produce economically better designs. This is especially important as underground mine projects are becoming more marginal. The difference between an initial design and an optimal one could be enough to make an undesirable project feasible. Additional contributions are regarding the optimization techniques presented. Both the stope geometry and stope sequence optimization methods provide a basis for future algorithm development. They can be adapted to specific projects for application purposes.

6.1 Stope Geometry Optimization

A method for optimizing initial stope designs has been presented. An initial design is modified such that its economic value is maximized. Economic value is calculated as the sum of all blocks from the expected economic model that are within the stope. Simulated annealing was used as the optimization algorithm. Given a stope that is represented as a

triangulated solid, the algorithm randomly cycles through the vertices shifting them in space. This is done until the annealing temperature reaches a user defined minimum. Simulated annealing was chosen for several reasons: it is a global optimization technique, the problem is of high dimensionality, and because of its ease of implementation.

It was shown that simulated annealing can improve on existing stope designs in different mining environments. A stope was designed for sublevel stoping and an improvement of 14% was made. One for cut-and-fill mining was also optimized yielding a gain of 32%. However, improvement for the cut-and-fill stope may be inaccurate due to a limitation of the algorithm regarding selection and moving of vertices. This will be explained in the future work section below.

The method for stope design optimization was made quite flexible. It can handle any triangulated solid; there are no limitations in terms of solids being convex. However, there are several requirements that must be met. An economic model distributed through a regular grid must be provided to optimize over. The triangulation of the stope design must be predefined. A fixed set of constraints was provided with which most realistic mining constraints can be accounted for. There may be other project specific constraints that cannot be handled.

6.2 Stope Sequence Optimization

This thesis has presented a framework for stope sequence optimization. The problem was kept simple in terms of underground mine scheduling; however, requirements for a more advanced sequencing algorithm have been discussed. The mining scenario considered involved a panel of stopes in a sublevel stoping environment. One crew was designated for development and stope preparation and another for ore extraction. Three constraints were applied regarding timing of events: (1) once a stope is chosen to be mined, it must be fully prepared and extracted before the next one in the sequence, (2) development can only take place towards one stope at a time, and (3) breaks in ore extraction are not permitted. Given this information, the optimization problem was to choose a sequential order to mine the stopes in that provides the maximum net present value. It is a combinatorial optimization problem.

Two optimization techniques were explored: simulated annealing and a logic driven algorithm. Simulated annealing was used because of its ease of implementation. Use of a logic driven algorithm was based on there being more to choosing an order than random perturbations. Sequence decisions were made based on information about the stopes and timing of operations. This information was used to calculate probabilities, which dictate the sequence. Both algorithms were able to improve the NPV from an initial sequence. With simulated annealing, the initial sequence was the first random selection. With PDM, the initial sequence was chosen greedily. Stope revenue was the basis for the greedy algorithm. Based on the results, simulated annealing found a better solution in less time.

As this was only a framework, there are limitations. Sequencing is very dependent on the mining method being used and attributes of the deposit. Geologic and geotechnical information control many of the constraints that come into sequencing. Therefore, developing an algorithm that can be applied to any mine is not realistic. However, a framework can be adapted to work in different environments. It provides a suitable starting point.

6.3 Future Work

Both stope geometry and stope sequence optimization could benefit from continued research. A wide variety of numerical optimization techniques that can be applied to these problems are available. Some may be more applicable to these problems than simulated annealing. In any case, other optimization techniques should be reviewed and tested on both problems. For stope geometry optimization, approximate gradient based methods such as simultaneous perturbation stochastic approximation (SPSA) could be attempted. Exact methods such as a systematic search may also be applicable as long as run-time does not become unrealistic.

There are many optimization algorithms that could be applied to stope sequencing. Simulated annealing that was applied performed random stope order exchanges. Exchange heuristics based on known information may be beneficial. Algorithms such as branch and bound can be formulated for stope sequencing. For branch and bound, an additional algorithm for calculating the bounds on sequence NPV would be required. Other optimization techniques such as genetic algorithms could be explored for stope sequence optimization.

Improvements are required for the probabilistic decision making algorithm. It was identified that as weights are updated, they eventually settle into a local minimum for each stope in the sequence. At this point, the sequence cannot be changed. Steepest descent was used for updating the weights. This could still be used, but in conjunction with techniques for escaping local minima such as tunneling. Another aspect of PDM that could be improved is the depth into the sequence the algorithm is searching. Currently, only one decision is made and then a NPV is calculated with the previous best order. It may be better to allow the algorithm to increase the number of decisions as the learning rate flattens. Once learning has halted for one decision at a time, try two, then three and so on.

Additional future work surrounds other aspects of the two problems. For stope geometry optimization, run-time when considering block refinement was substantially longer than with no refinement. Efficiency of block model-stope intersection needs improvement. Exact intersection of blocks with stope solids should be considered as well. Another requirement was identified when optimizing the cut-and-fill stope. Currently one vertex is shifted at a time, but for stopes such as that for the cut-and-fill scenario it would be beneficial if subsets of vertices could be shifted in tandem. This would allow a stope's cross-section to remain fixed for example.

Another area of interest for stope geometry optimization is automatic insertion and deletion of vertices as optimization proceeds. The triangulated mesh could undergo adaptive refinement or coarsening. For example, refinement could be used to increase accuracy in areas where faces grow to a point where selectivity of equipment is underrepresented. For cut-and-fill, as parts of a drift are shifted, more vertices will be needed to accurately represent corners and curves. On the other hand, vertices that converge very close together during optimization could be merged.

Improvements in terms of constraints can also be made. Constraints sometimes depend on the type of rock being mined and other geological features such as faults and joints. To incorporate these constraints, stope modification would take place over other models such as geology or geotechnical models as well as the economic model. The economic model that optimization is based on is also limited. Many parameters can be incorporated into the economic model; however, those that cannot should be incorporated into stope geometry optimization. One of these is the cost for installing rock support, which is dependent on the geometry of the stope and surrounding geological features.

For stope sequence optimization, the list of future work items would be extensive. Generally, future work consists of adding complexity to the framework to represent more realistic mining scenarios. New data structures could be incorporated to represent geometric components of a mine including stopes, drifts, ramps, and crosscuts. More information about equipment that is being used including its operating parameters and costs would improve accuracy of results. A wide variety of other constraints exists for various mining projects. Many are related to safety and limit when and where mining can take place. Others relate to meeting production requirements and ore blending. Additional future work is to apply different optimization techniques with the existing framework and as the scenario is made more complex.

References

1. Aarts E. and Korst J., Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing, John Wiley & Sons Ltd., 1989

2. Boyd S. and Vandenberghe L., Convex Optimizations, Cambridge University Press, 2004

3. Brusco M.J. and Stahl S., Branch-and-Bound Applications in Combinatorial Data Analysis, Springer Science + Business Media Inc., 2005

4. Coley D.A., An Introduction to Genetic Algorithms for Scientists and Engineers, World Scientific Publishing Co., 1999

5. Coppin B., Artificial Intelligence Illuminated, Jones and Bartlett Publishers, 2004

6. Cormen T.H., Leiserson C.E., Rivest R.L., and Stein C., Introduction to Algorithms, Second Edition, Massachusetts Institute of Technology, 2001

7. Gambardella L.M. and Dorigo M., An Ant Colony System Hybridozed with a New Local Search for the Sequential Ordering Problem, INFORMS Journal on Computing, Vol. 12, No. 3, 2000

8. Hand D., Mannila H. and Smyth P., Principles of Data Mining, The MIT Press, 2001

9. Hartman H.L. Editor, SME Mining Engineering Handbook, Society for Mining, Metallurgy, and Exploration, Inc., Littleton, Colorado, 1992.

10. Hartman H.L. and Mutmansky J.M., Introductory Mining Engineering, John Wiley & Sons, Inc., 2002

11. Hearn D. and Baker M.P., Computer Graphics with OpenGL, Pearson Prentice Hall, 2004

12. Heiser W.J. and Groenen P.J.F., The Tunneling Method for Global Optimization in Multidimensional Scaling, Psychometrika, Vol. 61, No. 3, 529-550, September 1996

13. Hoek E., Practical Rock Engineering, Course notes, 2000 edition

14. Hutchinson D.J. and Diederichs M.S., Cablebolting in Underground Mines, Bitech Publishers Ltd., 1996

15. Jones T., AI Application Programming, Charles River Media, 2003.

16. Neapolitan R. and Naimipour K., Foundations of Algorithms Using C++, Jones and Bartlett Publishers, 2004

17. Papadimitriou C.H., and Steiglitz K., Combinatorial Optimization Algorithms